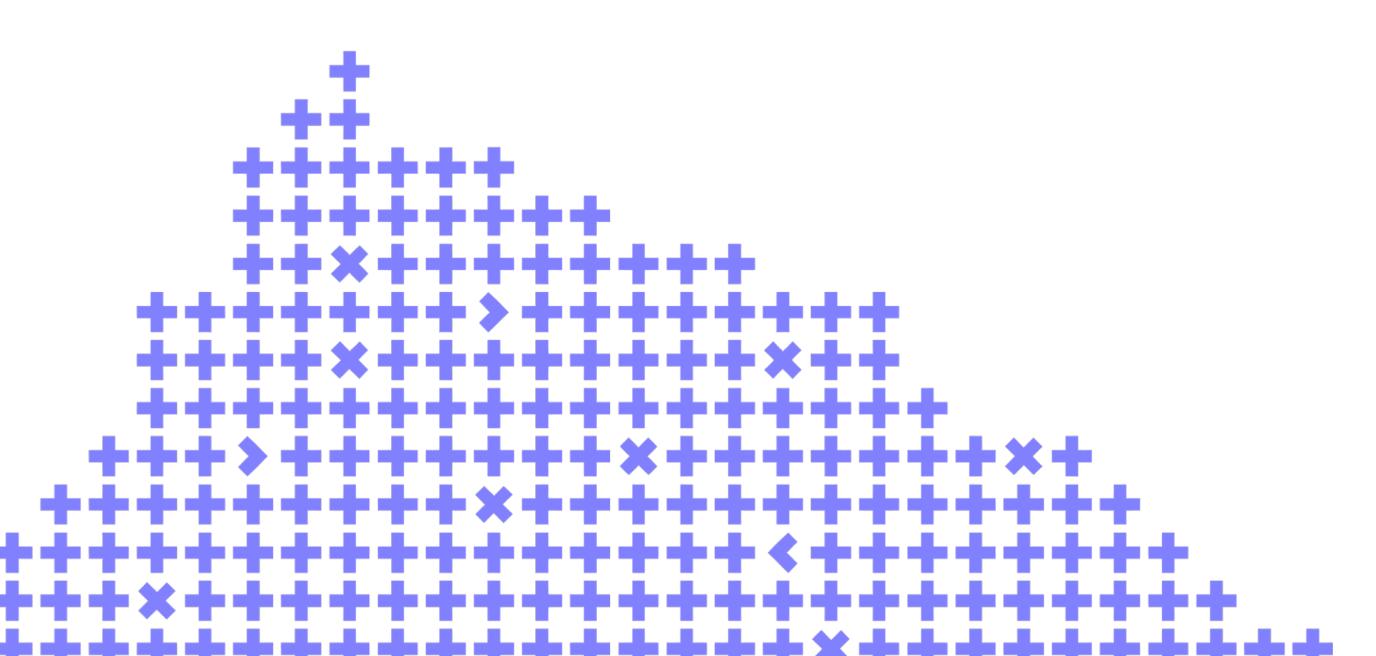
# How we reduced logs costs by moving from Elasticsearch to Grafana Loki

Igor Latkin





Co-organizer



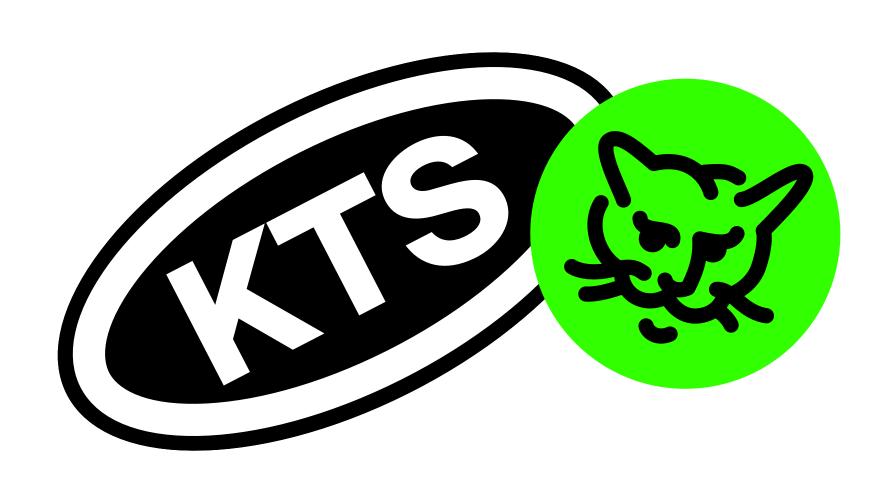
#### MHO WW I;



#### **IGOR LATKIN**

Co-founder & System Architect @ KTS

- Corporate systems
- Non-standard projects
- Mobile
- DevOps





#### AGENDA



- 1. Log collection task
- 2. Loki architecture
- 3. Our journey of logs transferring from ES to Loki
- 4. (bonus) Loki configuration tips & tricks

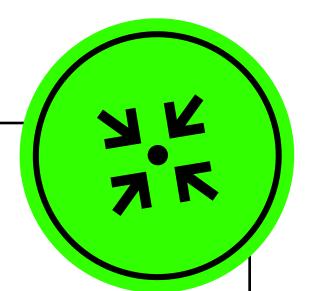
#### WHAT THIS TALK IS NOT



- Loki or Elasticsearch tutorial
- Loki vs Elasticsearch comprehensive comparison
- Complete set of instructions how to transfer logs in your environment

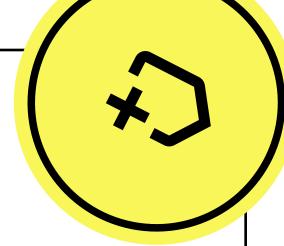
#### LOG COLLECTION PROBLEM





How to collect

How to store



How to extract metadata

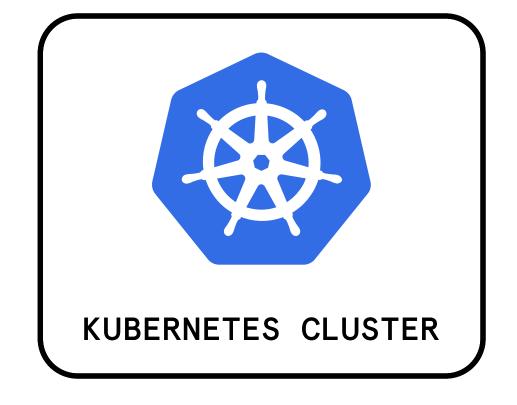


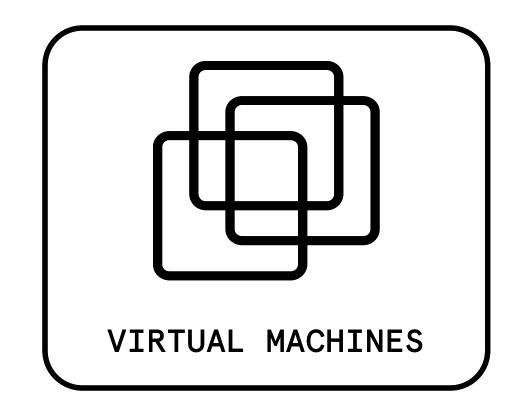
to query

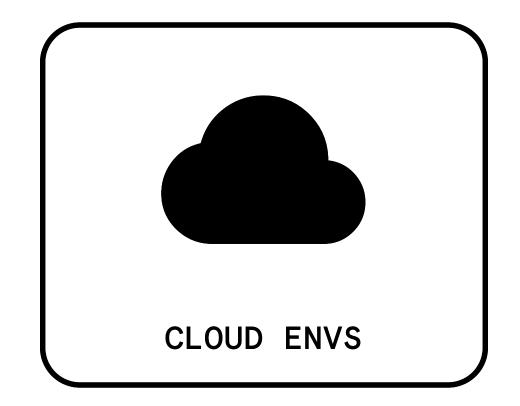


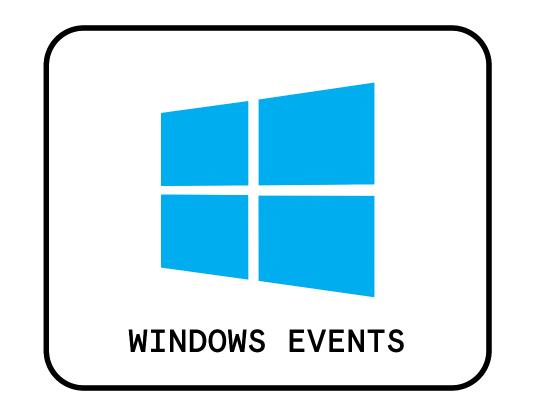
#### MULTIPLE SOURCES





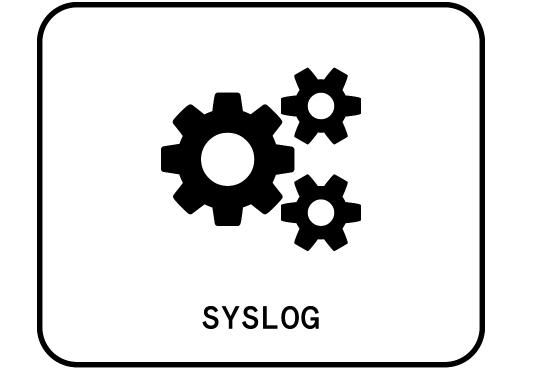


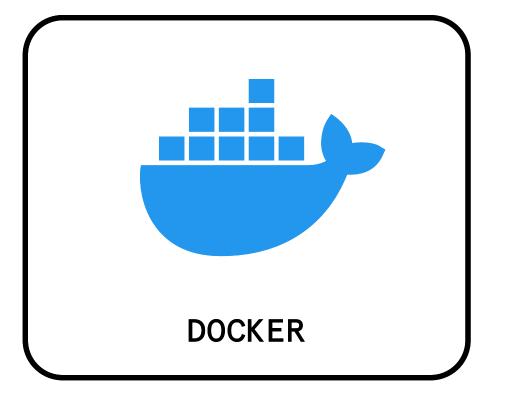






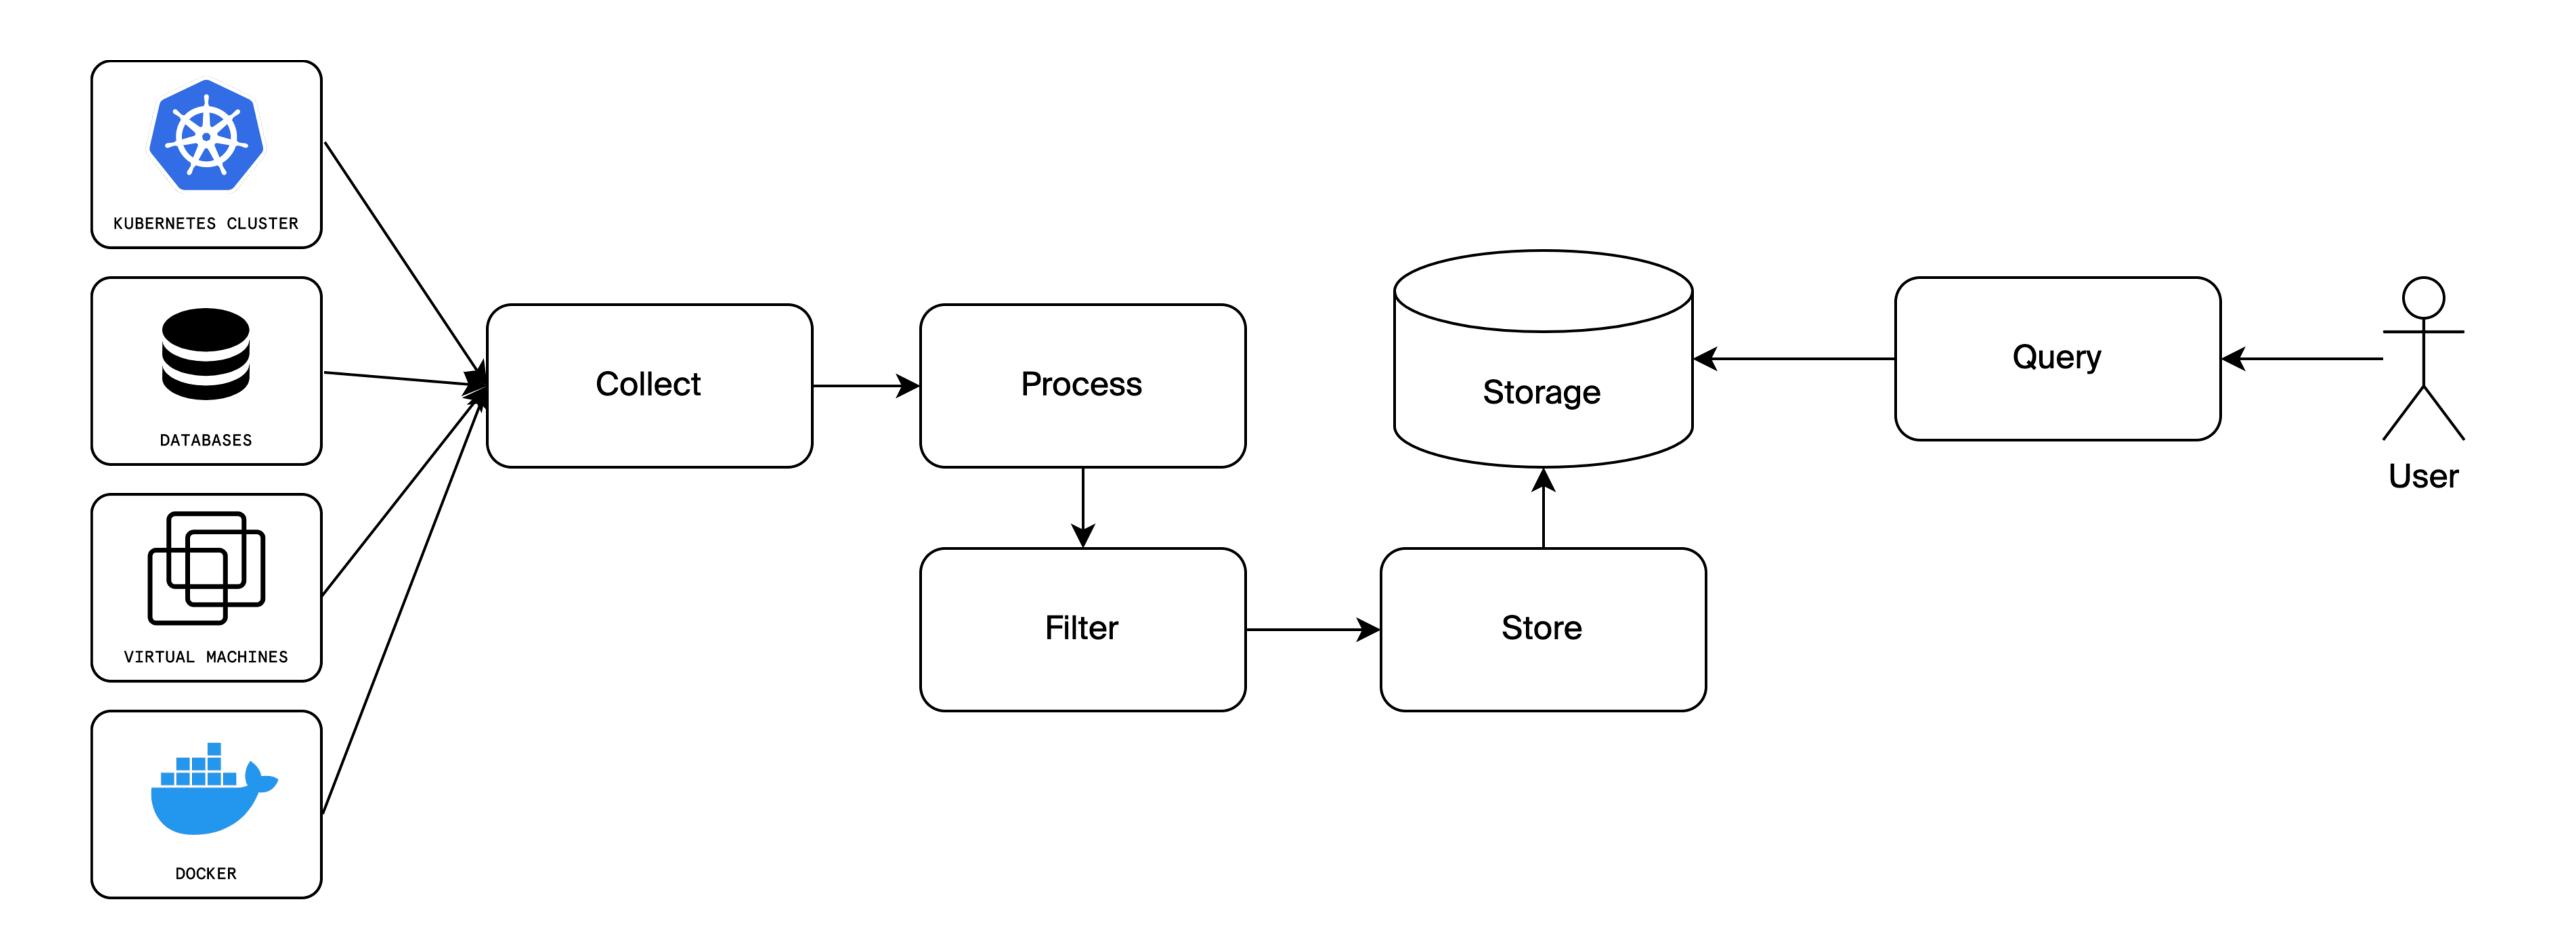






#### LOGGING PIPELINE





#### WE ALL KNOW THEM



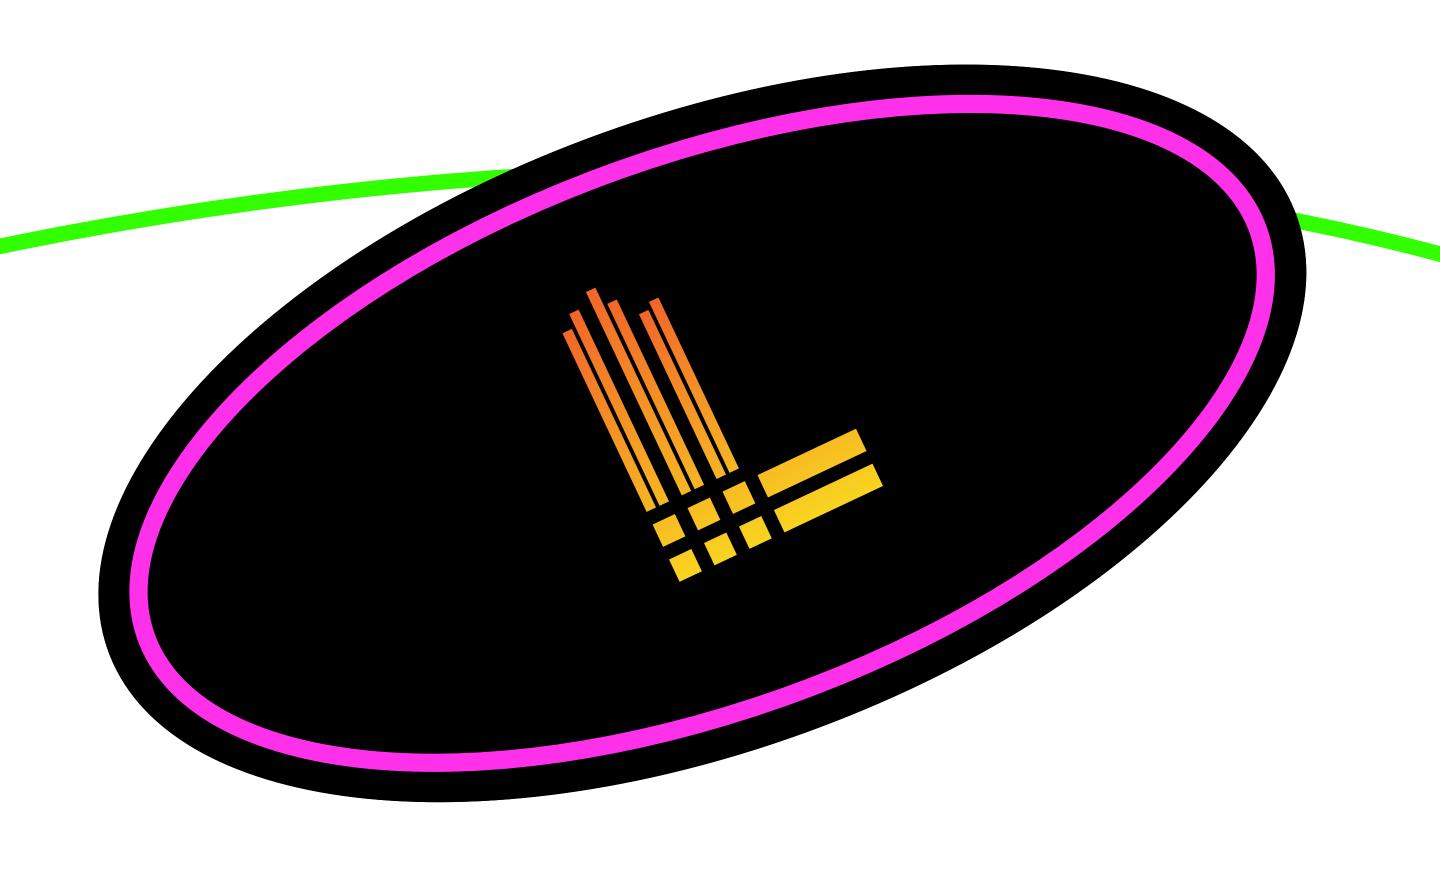






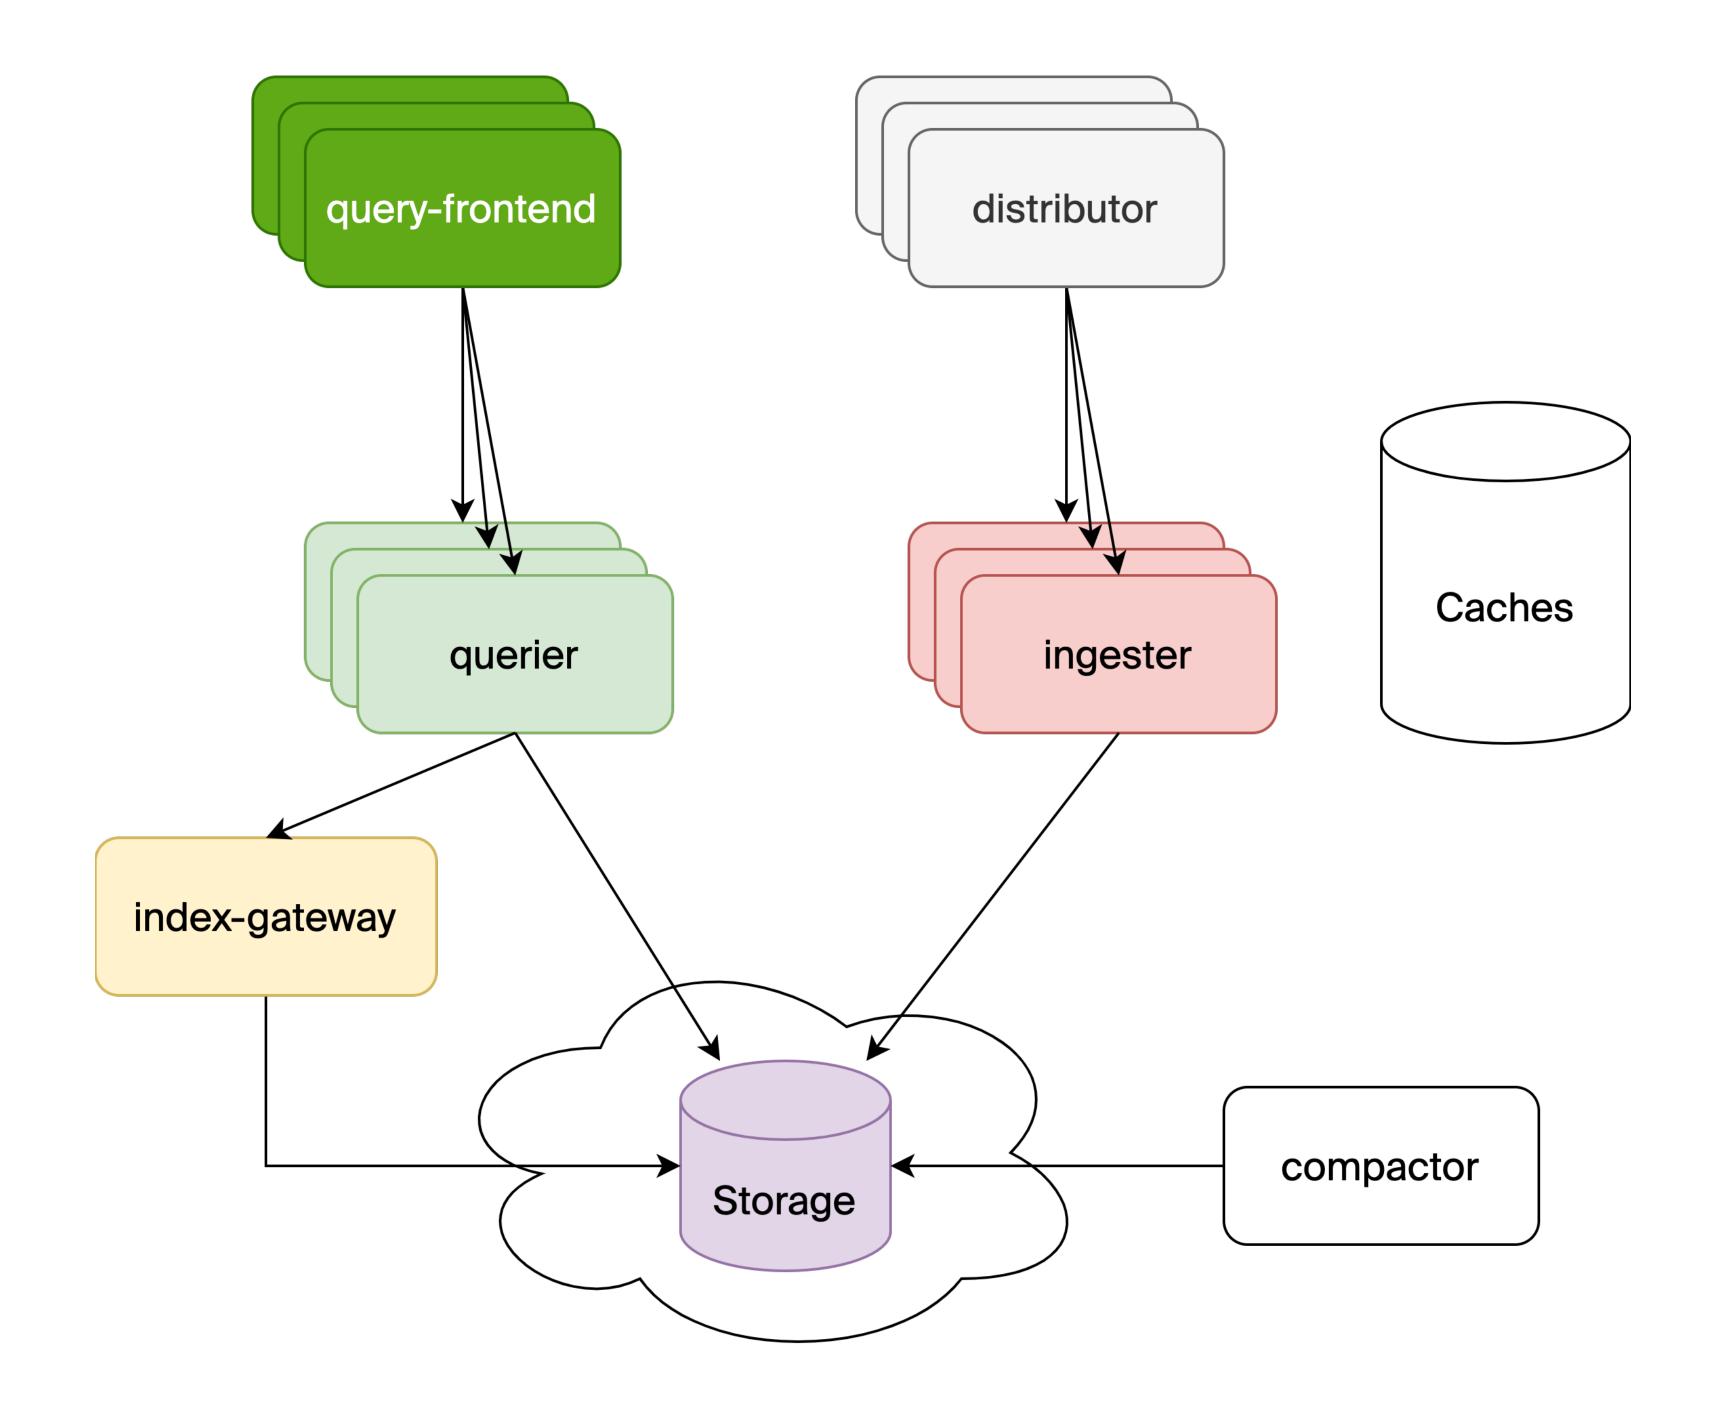






### -> LOKI ARCHITECTURE





#### INDEXING



2019-12-11T10:01:02.123456789Z {app="nginx",cluster="us-west1"} GET /about

**Timestamp** 

with nanosecond precision

**Prometheus-style Labels** 

key-value pairs

Content

logline

indexed

unindexed

#### STREAMS

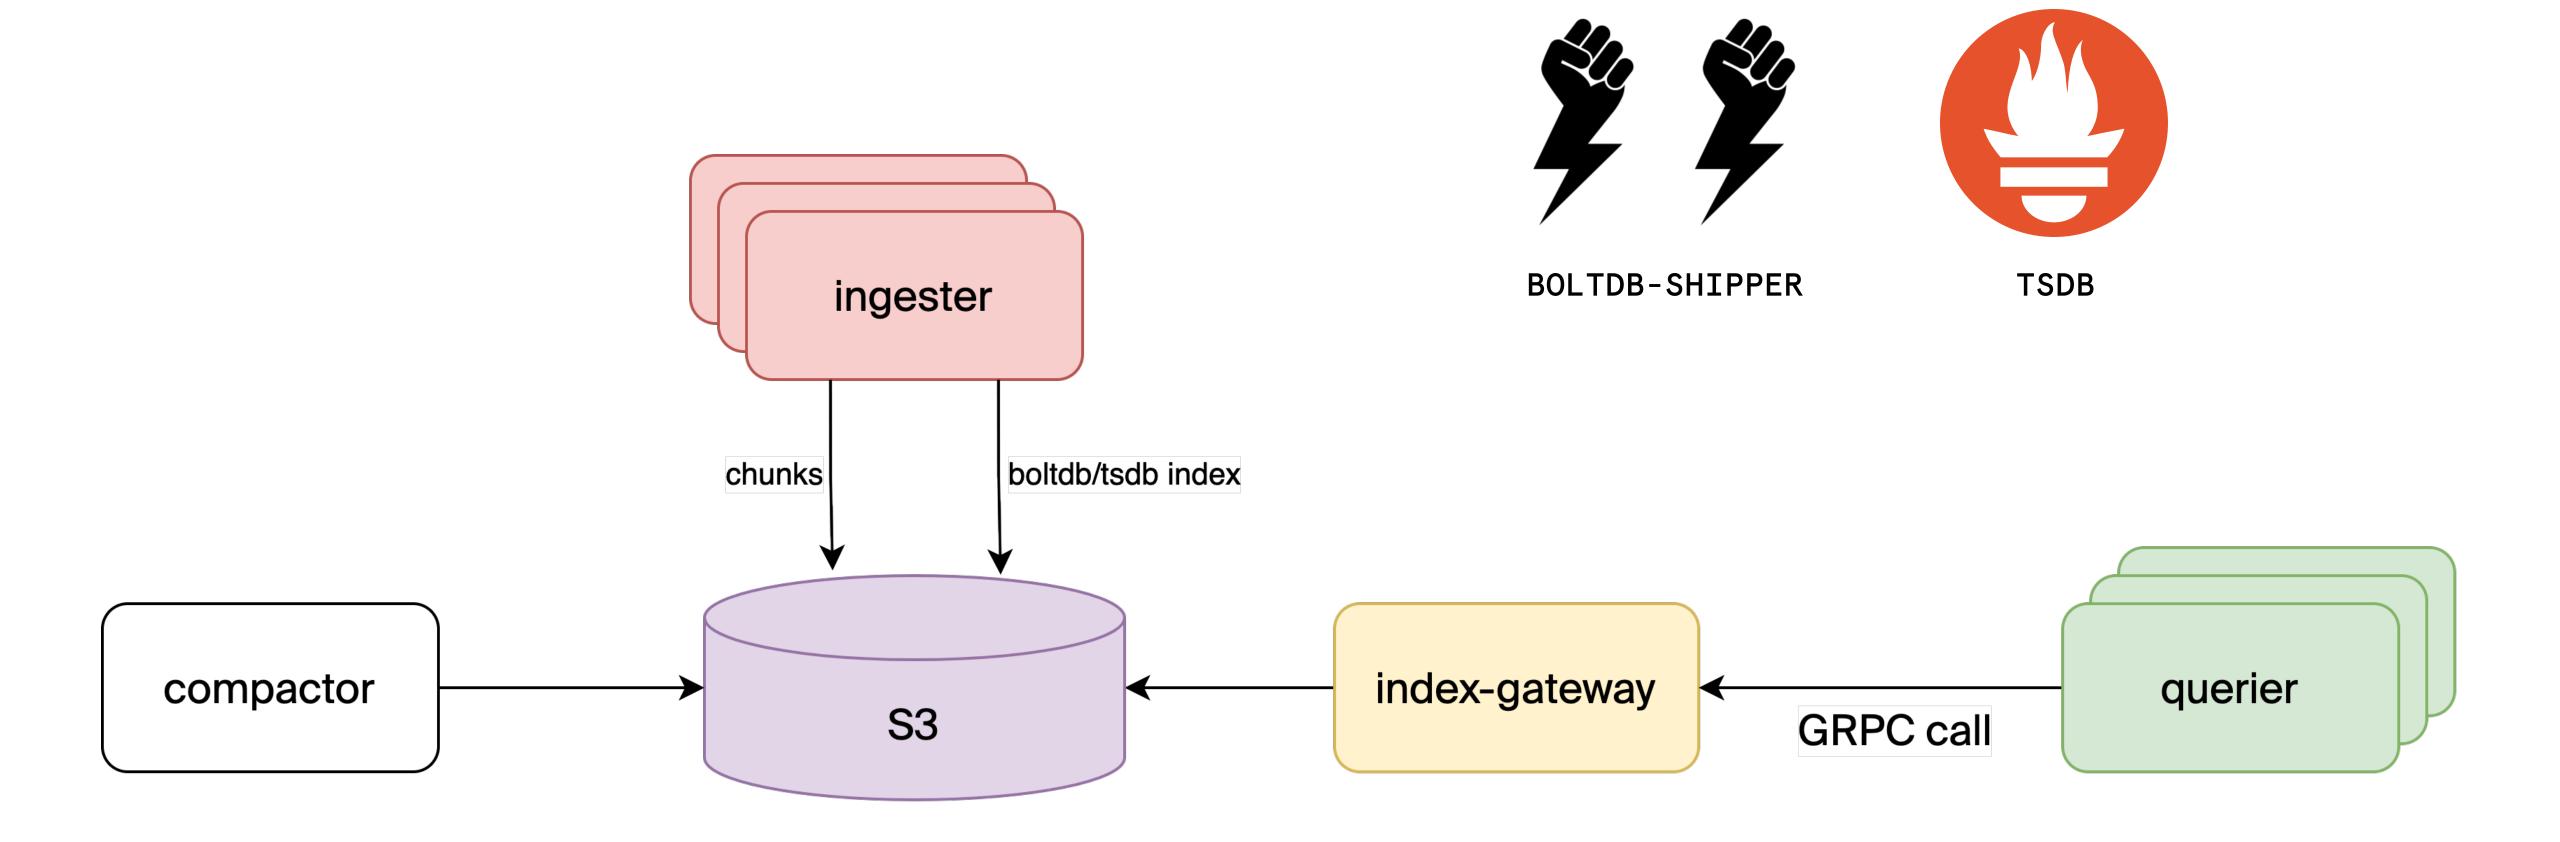


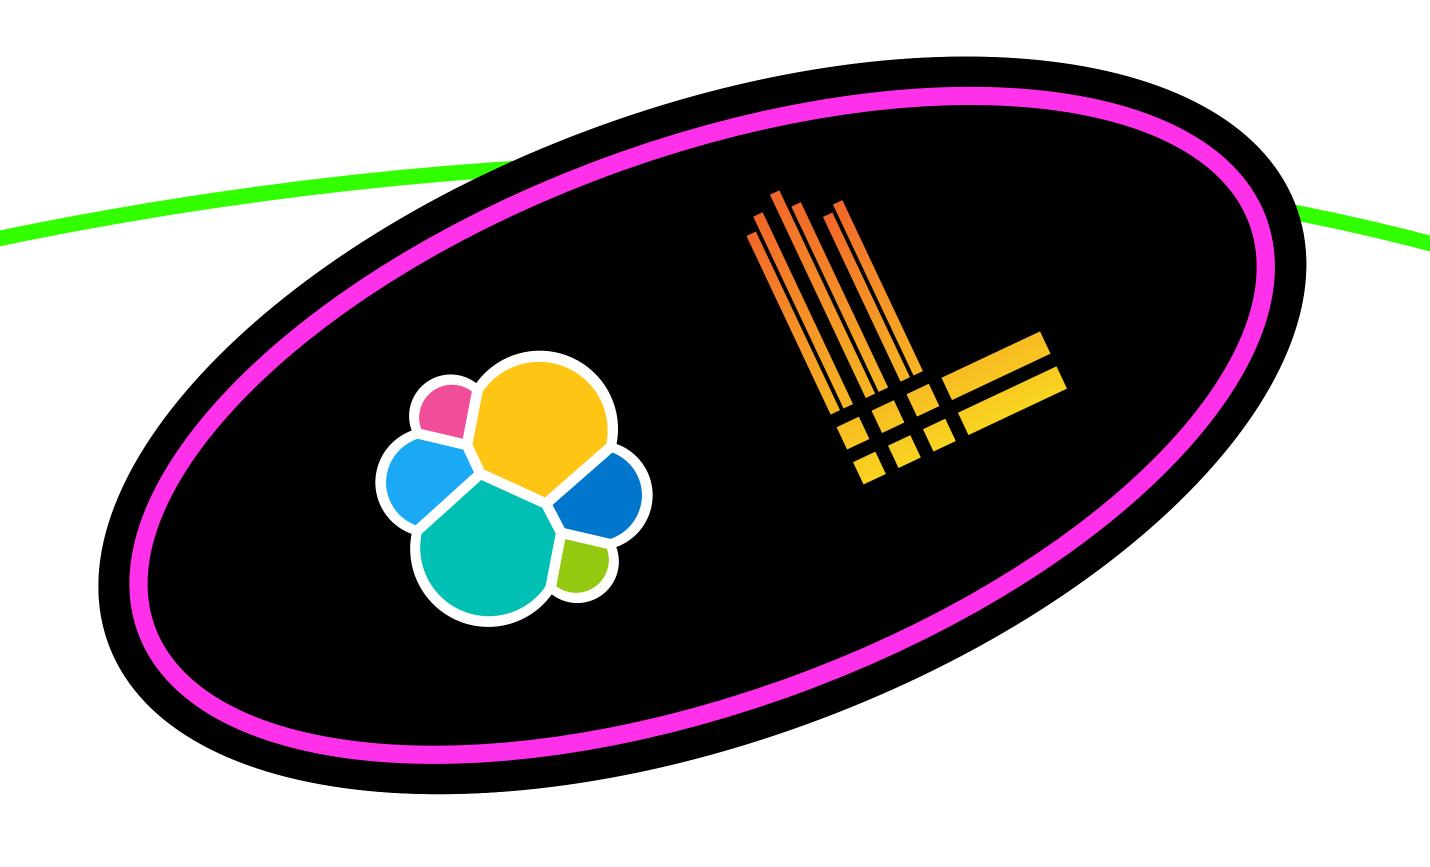
Logs		
Labels	Message	
<pre>{component="printer", location="f2c16", level="error"}</pre>	"Out of paper"	chunk #1 stream 1
<pre>{component="printer", location="f2c16", level="error"}</pre>	"Too much paper"	Sucami
<pre>{component="supplier", location="f2c16", level="info"}</pre>	"Paper exhasted"	chunk #2 stream 2

#### CHUNKS & INDEX

#### CHUNKS & INDEX > S3

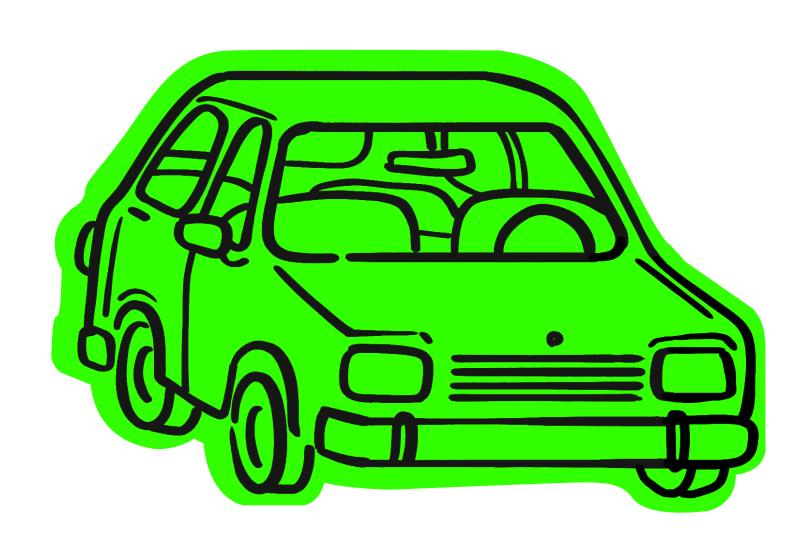






## → ELASTICSEARCH → LOKI/

#### LO WOAE S MHA ME DECIDED



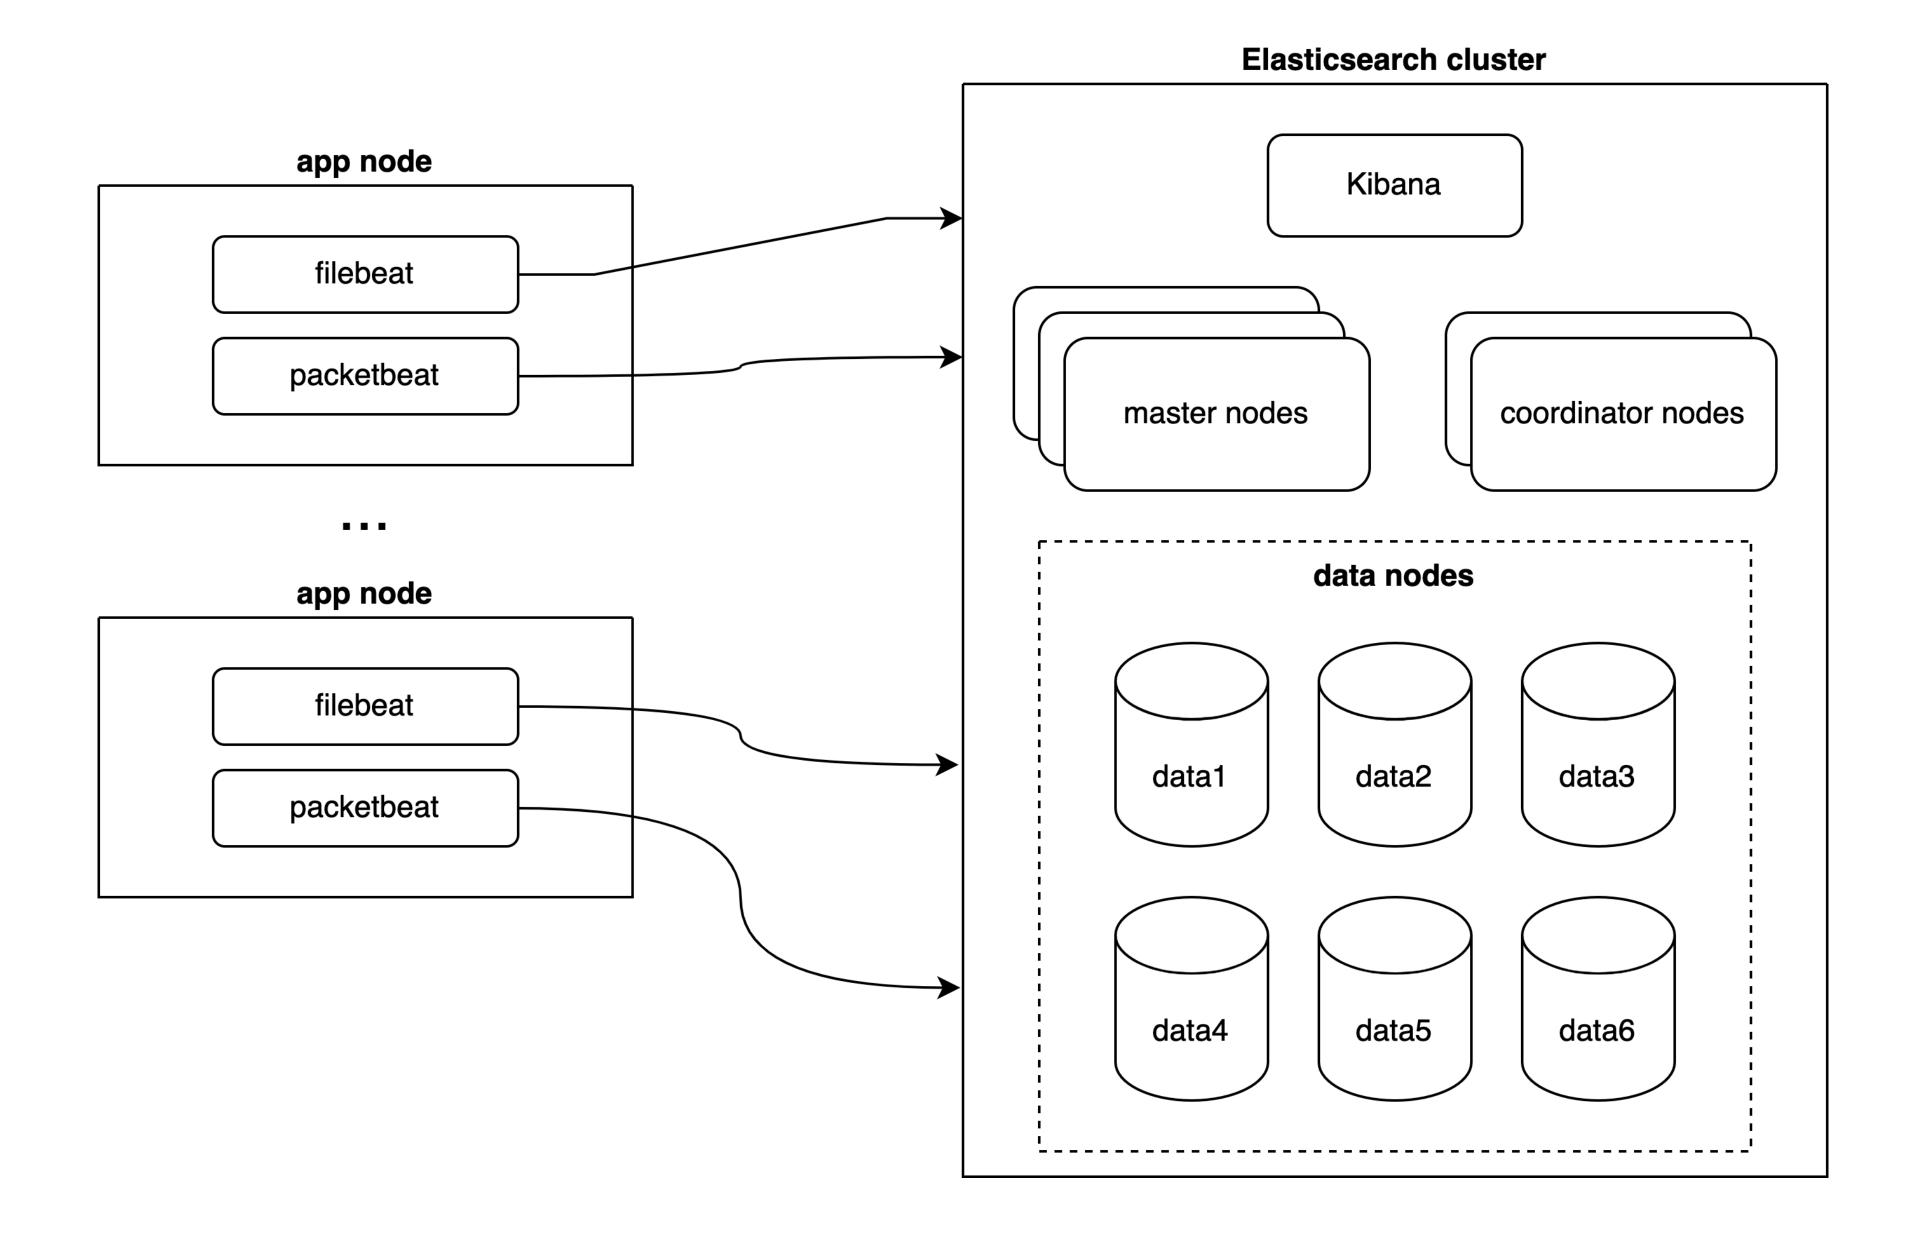
#### MOTIVATION TO DITCH ELASTIC



- Huge money burden
- 2 Hard maintainability

#### ELASTICSEARCH-BASED INFRA





#### MIGRATION PLAN



- 1. Deploy Loki cluster into Kubernetes
- 2. Deploy S3 storage
- 3. Start collecting live logs in Loki
- 4. Solve troubles, constantly reconfigure Loki to handle the load
- 5. Transfer old logs
- 6. Fail multiple times
- 7. Succeed, wait for the transfer to complete
- 8. Demolish Elasticsearch cluster

#### MIGRATION PLAN

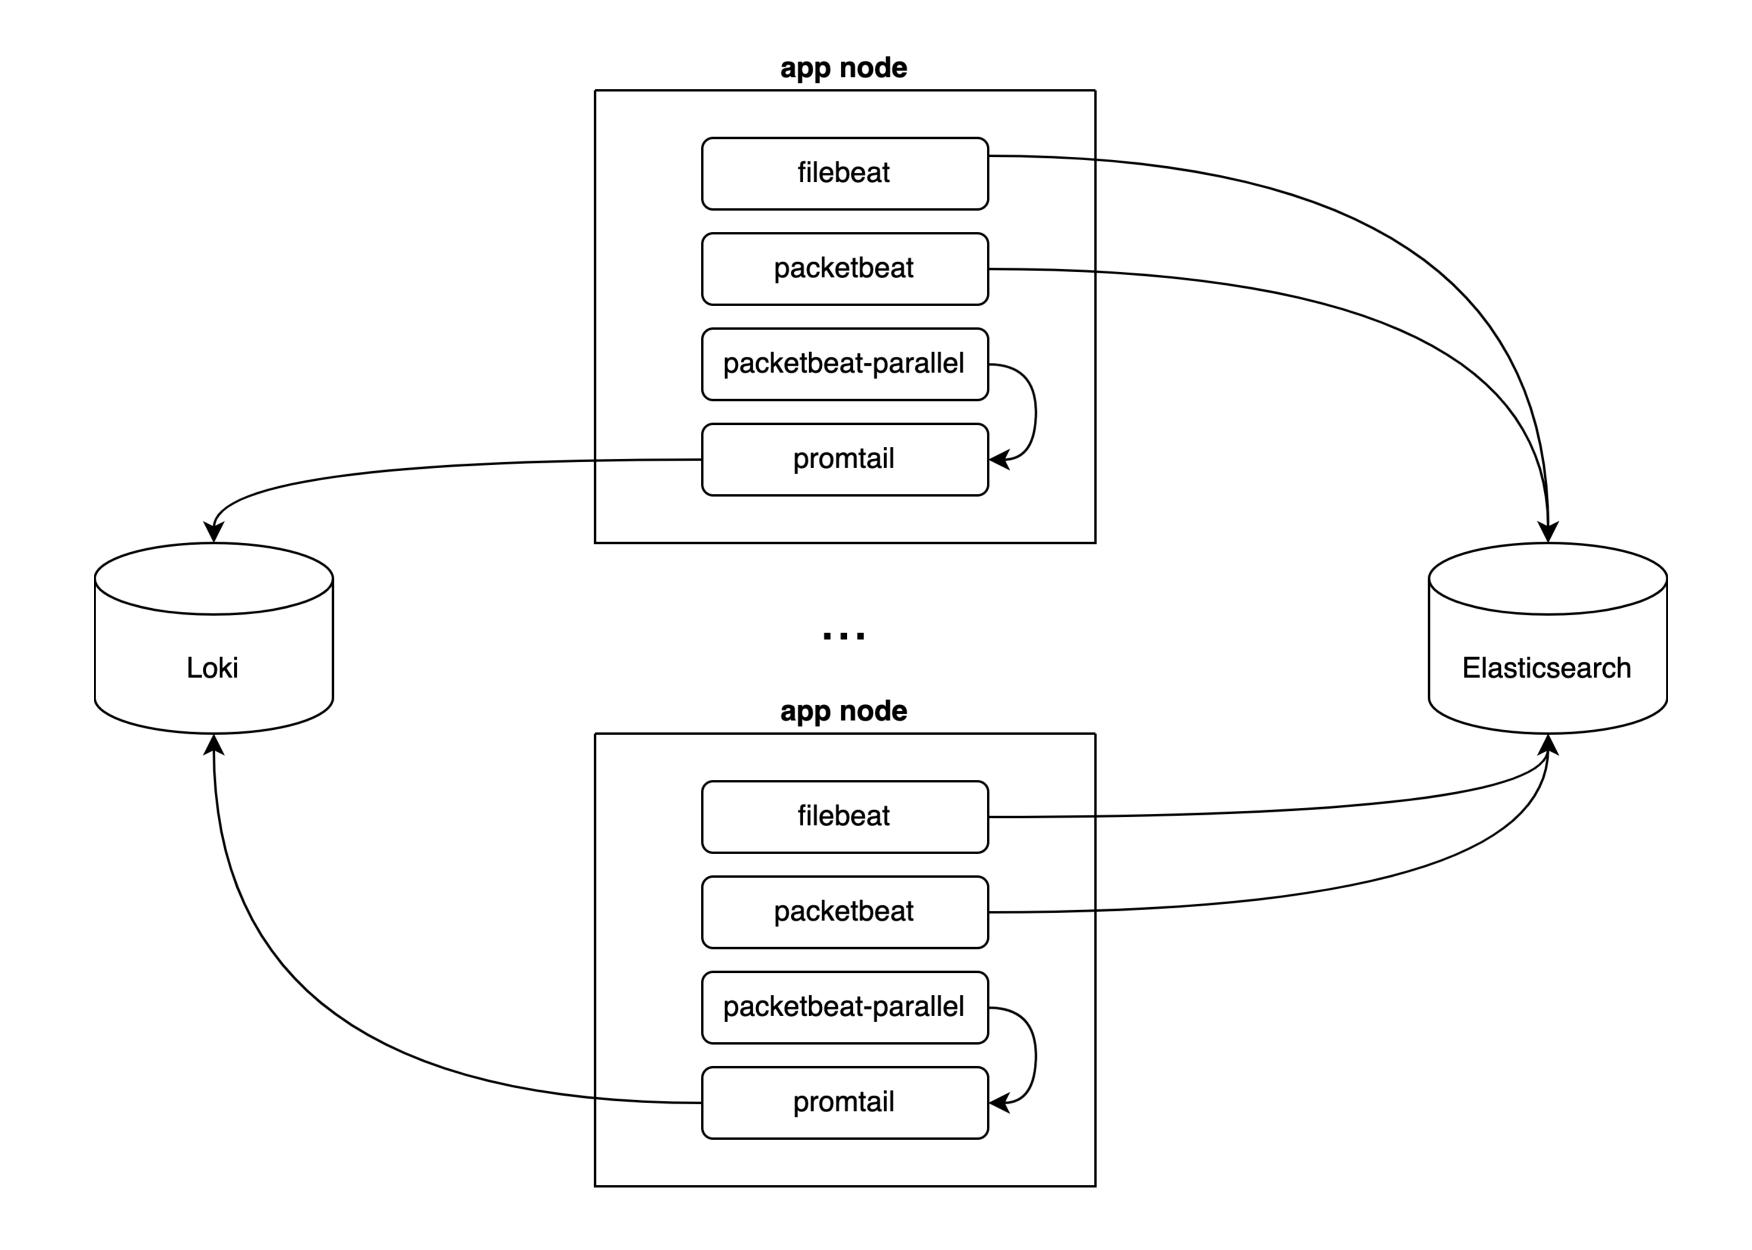


- 1. Deploy Loki cluster into Kubernetes
- 2. Start collecting live logs in Loki
- 3. Transfer old logs

# COLLECTING LIVE LOGS IN LOKI

#### NOT BREAKING EVERYTHING





#### COLLECTING LOGS FROM FILES



- Application logs
- Packetbeat logs
- Postgres logs (new)
- → 1C logs

```
1 scrape_configs:
 2 - job_name: system
     static_configs:
     - targets:
         - localhost
      labels:
         job: omni_services
         __path__: /var/log/omni/services/*/*log
10
11
     - targets:
         - localhost
12
      labels:
13
14
        job: packetbeat
15
         __path__: {{ promtail_packetbeat_folder }}/output
16
     - targets:
17
         - localhost
18
      labels:
19
20
         job: postgres_logging
         __path__: {{ promtail_postgres_folder }}/*/log/*log
21
```

#### EXTRACTED LABELS



app (50)	env (1)	job (3)	level (5)	node_name (4)
camunda	prod	1C_logging	DEBUG	1c-prod-2
catalog		omni_services	ERROR	prod-db01
cession-service		postgres_logging	INFO	prod-linux-1
customer-checker			TRACE	prod-linux-2
customer-data			WARN	
digital-subscriptions				
discount-service				
//	//	//	//	//

domain (64)	http_method (7)	http_status (19)	network_direction (2)	status (2)
	CONNECT	200	egress	Error
	DELETE	204	ingress	OK
	GET	301		
	HEAD	302		
	OPTIONS	303		
	POST	304		
	PUT //	400	/,	//

#### TRANSFER OLD LOGS

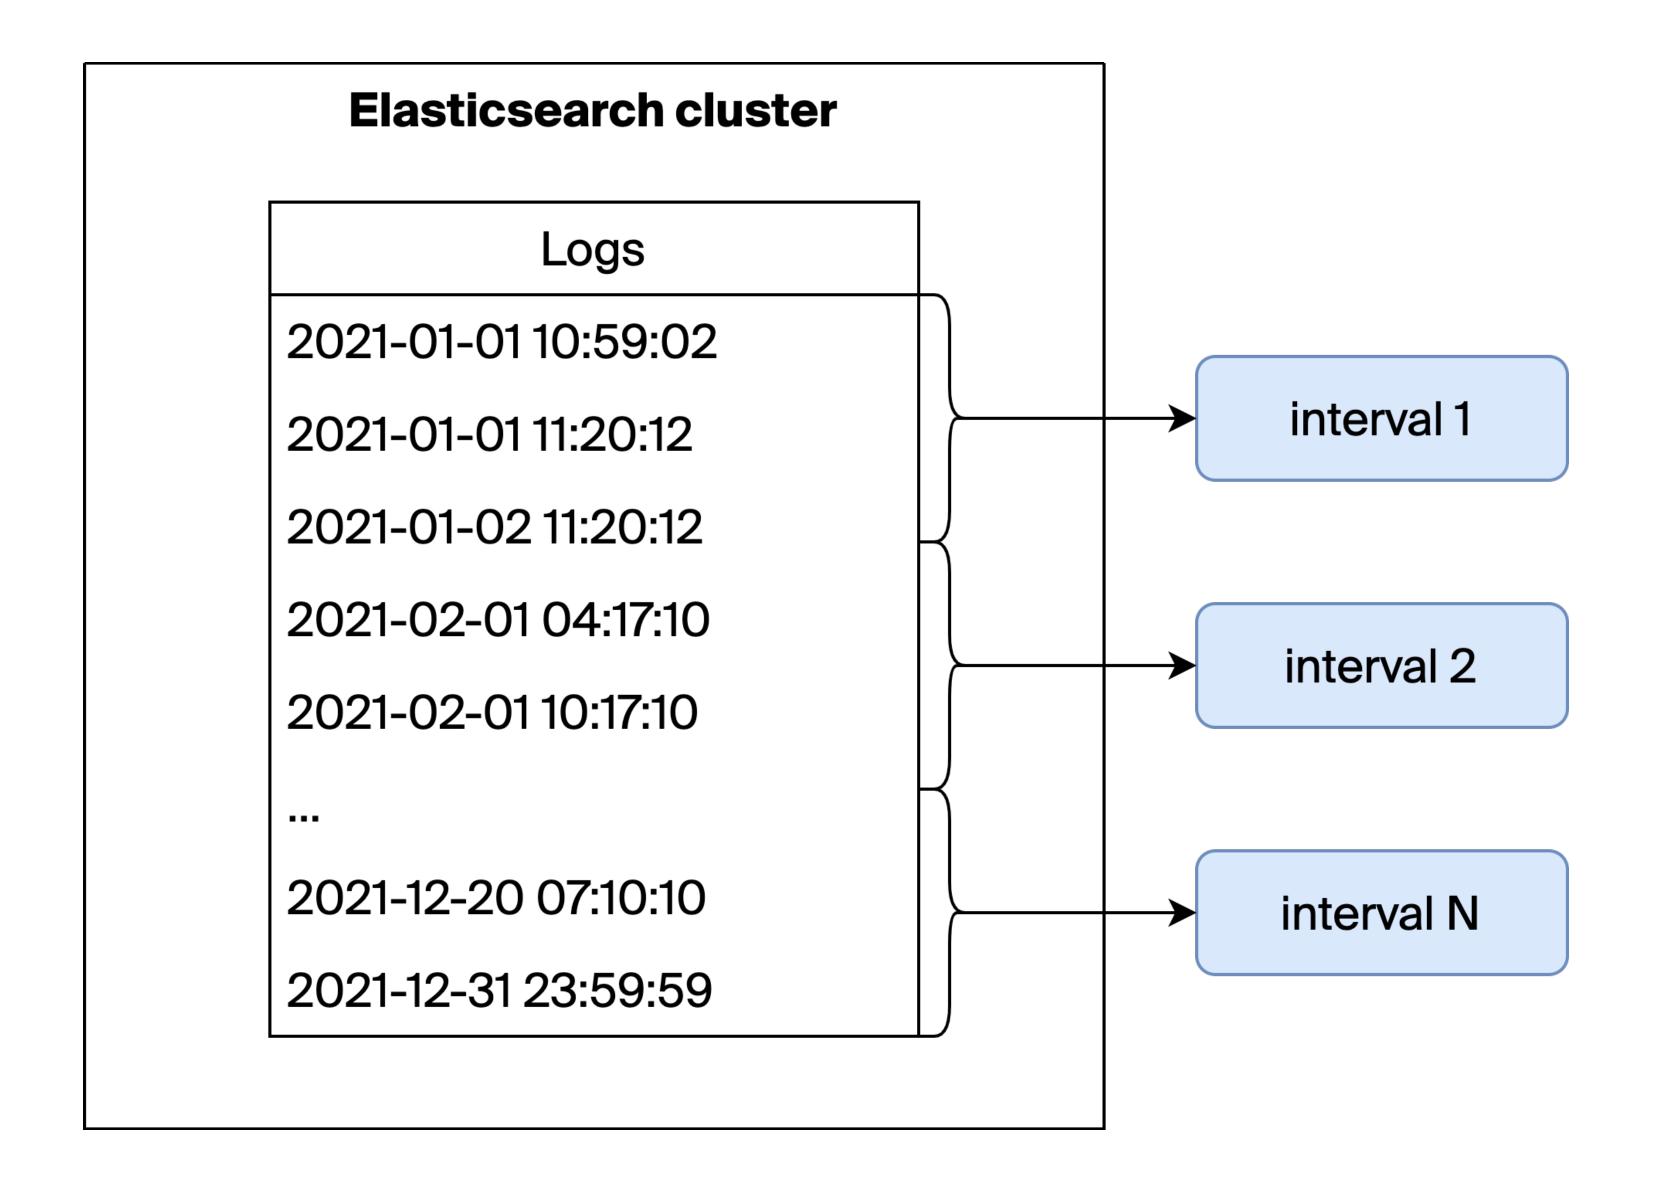
#### WHAT CONCERNED US?



- →1 year of log data
- -> 25 TB of storage

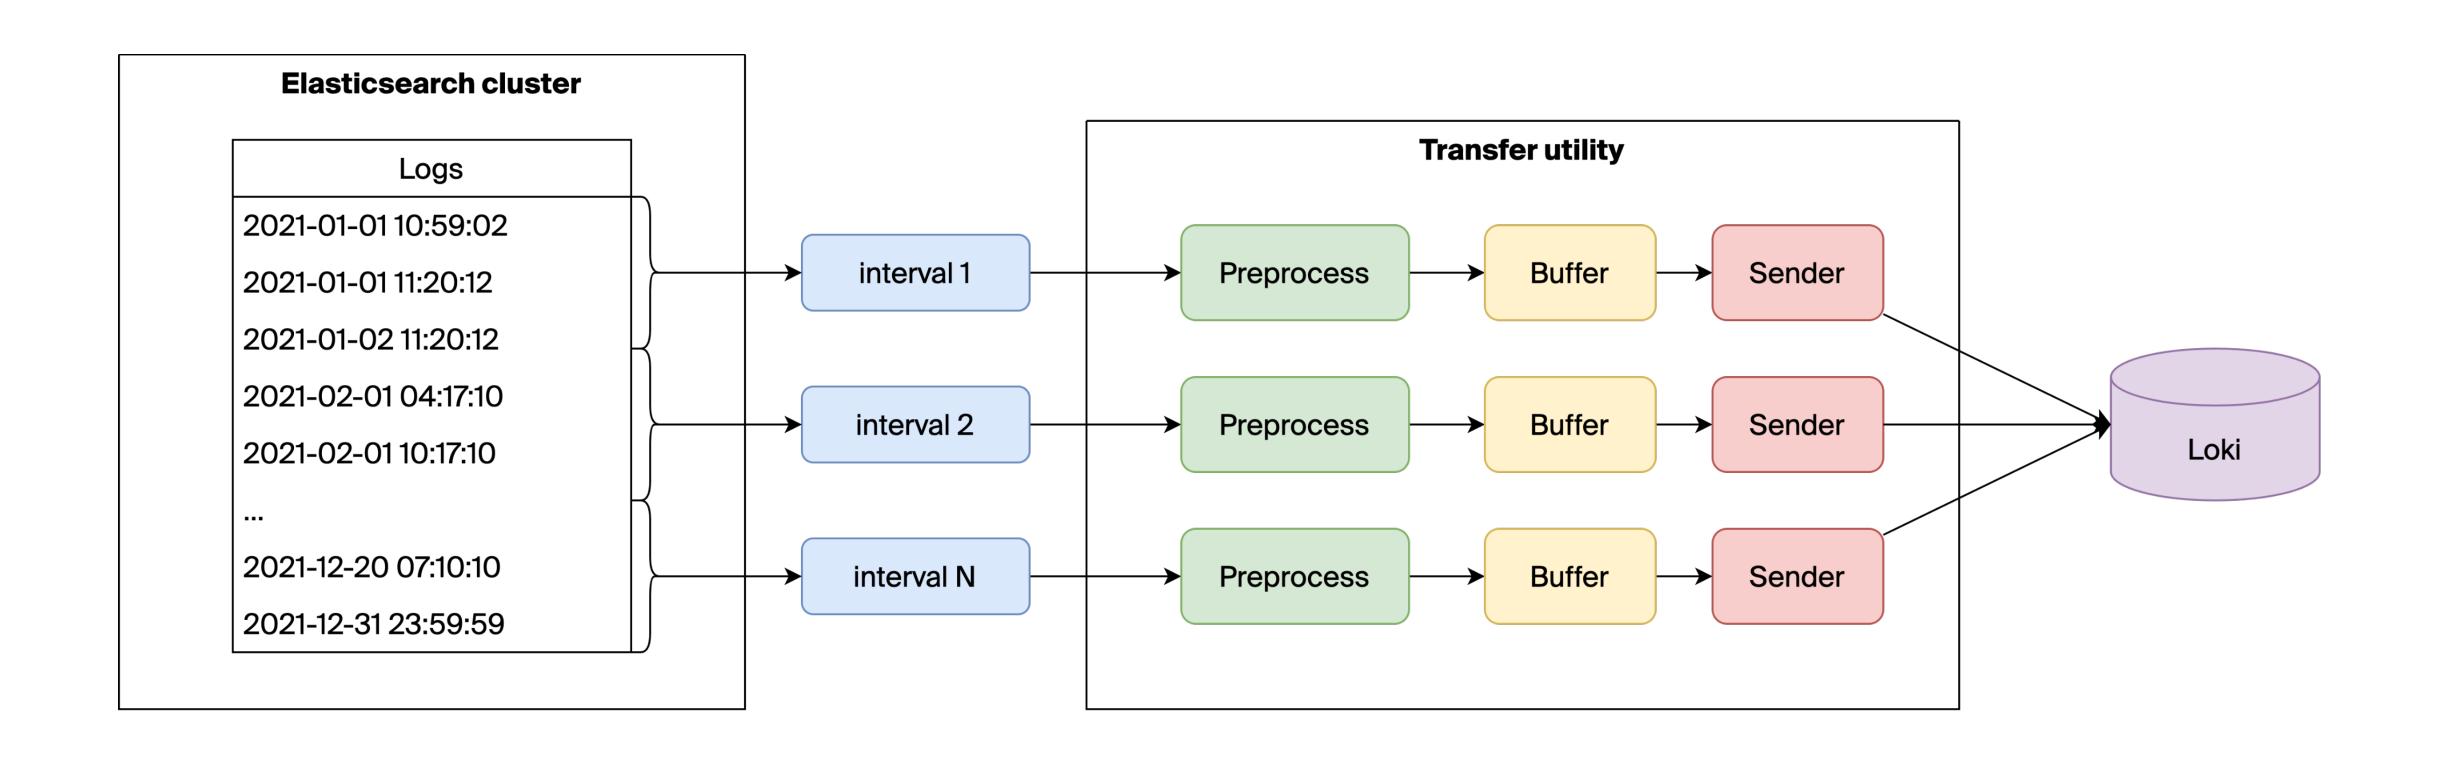
#### INITIAL ASSUMPTION





#### INITIAL ASSUMPTION





#### LOKI CONFIG



```
1 limits_config:
    reject_old_samples: false
    ingestion_rate_mb: 100
    ingestion_burst_size_mb: 30
    per_stream_rate_limit: "150MB"
    max_streams_per_user: 0
    max_global_streams_per_user: 0
10
    retention_period: 8928h
```

#### TESTING OUR ASSUMPTION



entry too far behind, oldest acceptable timestamp is: 2021-01-01T01:00:00

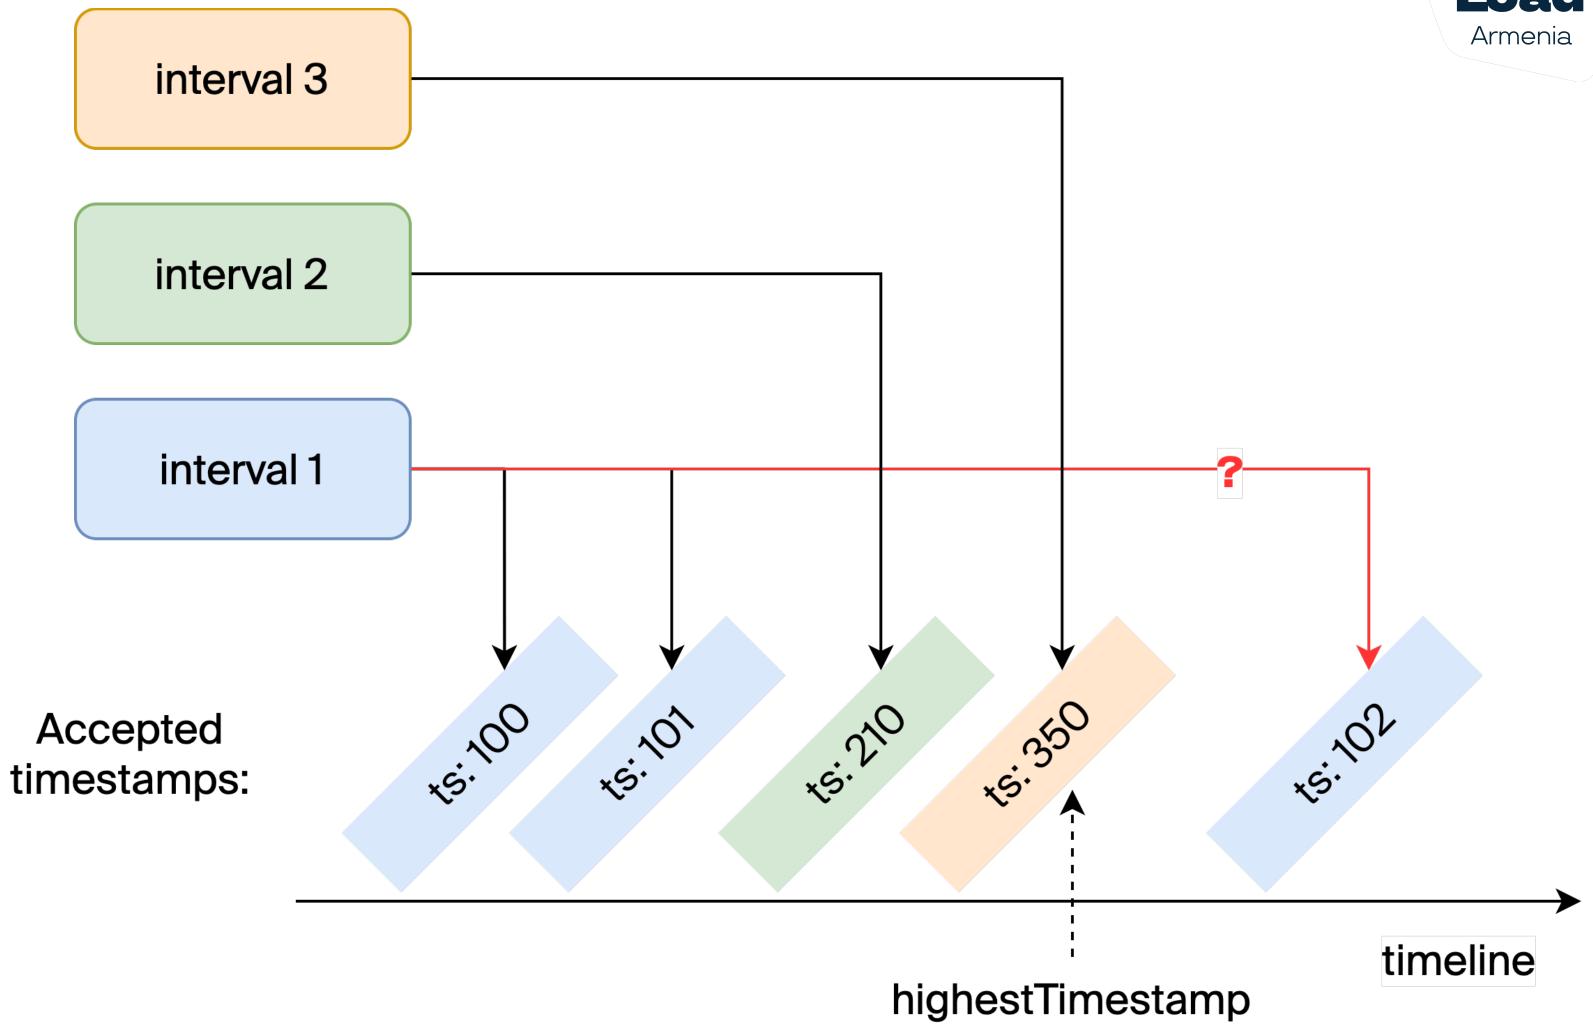
#### WHY IS THAT?



This is the error

#### MIND EXPERIMENT





Minimal accepted timestamp is:

highestTimestamp - maxChunkAge / 2

#### TWEAK CONFIGURATION

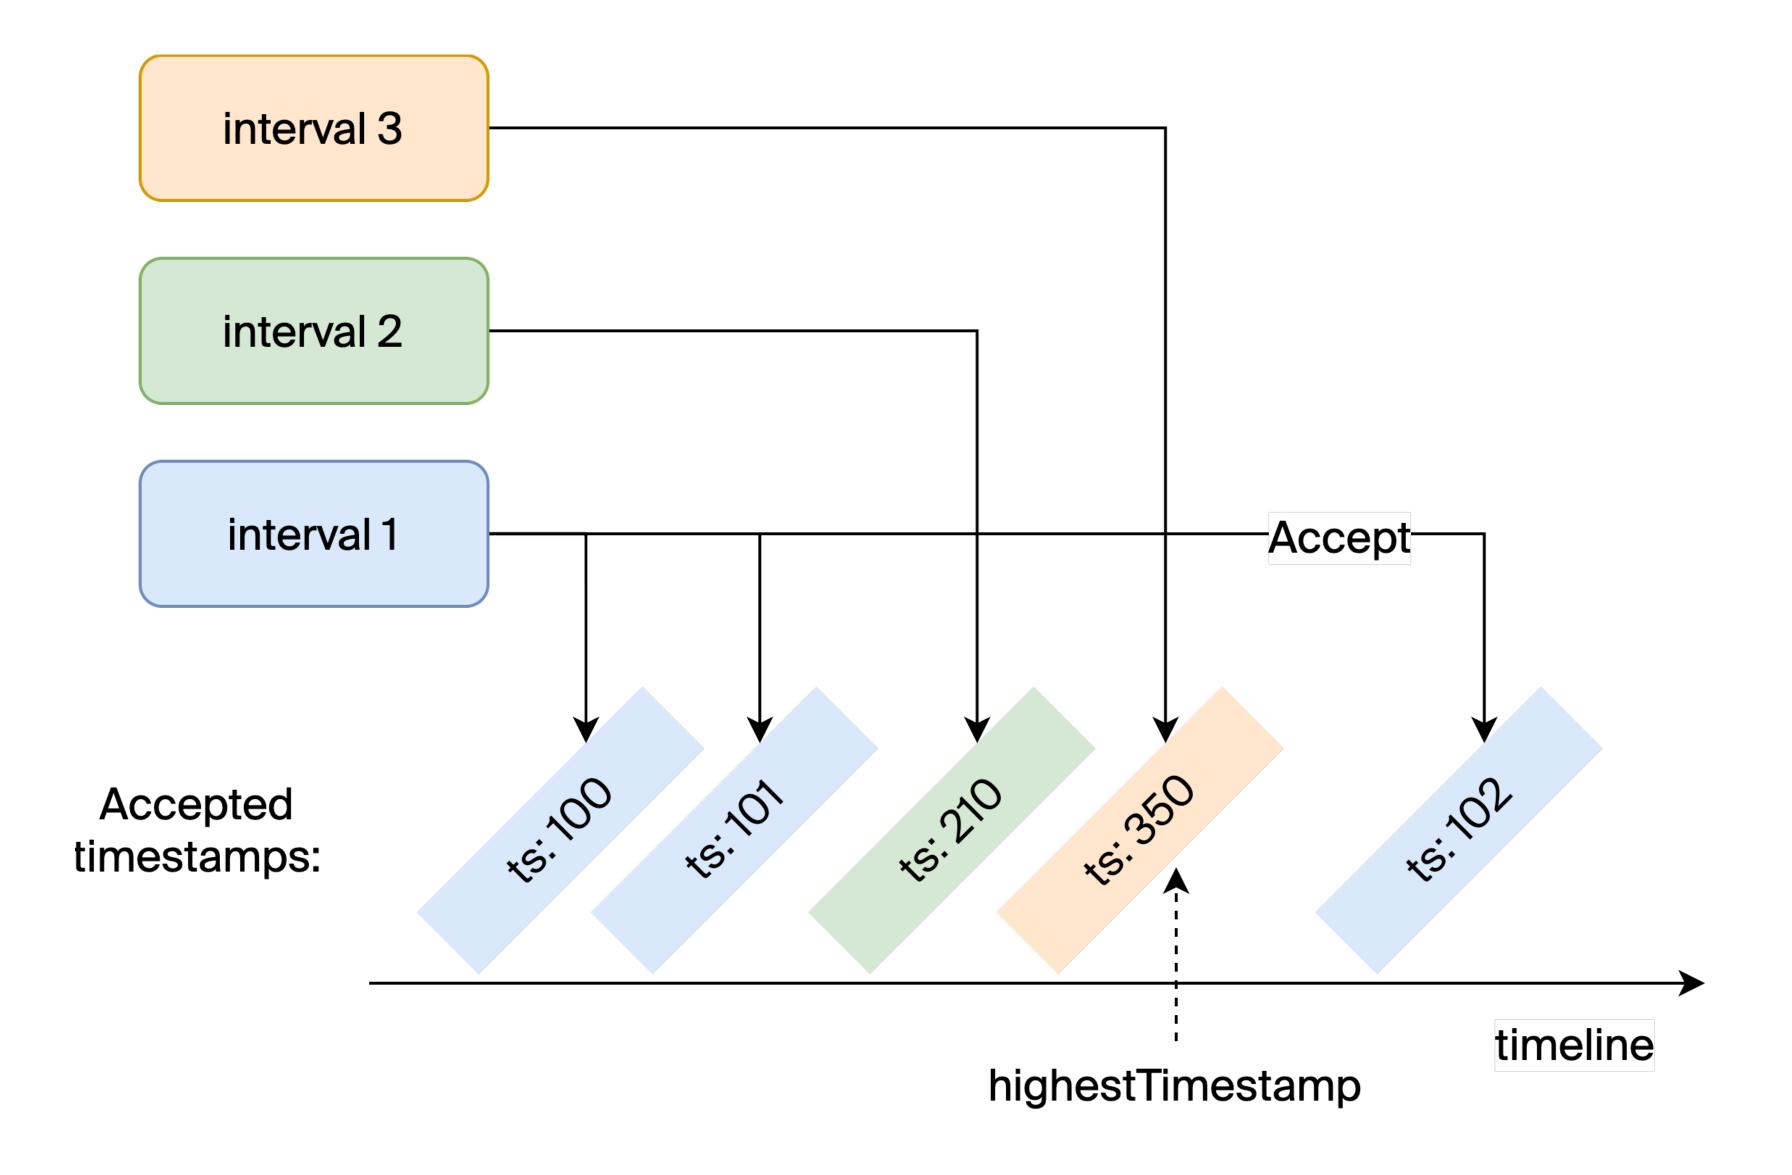


```
1 ingester:
```

2 max\_chunk\_age: 8760h # 365d

#### SEEMS TO BE FIXED





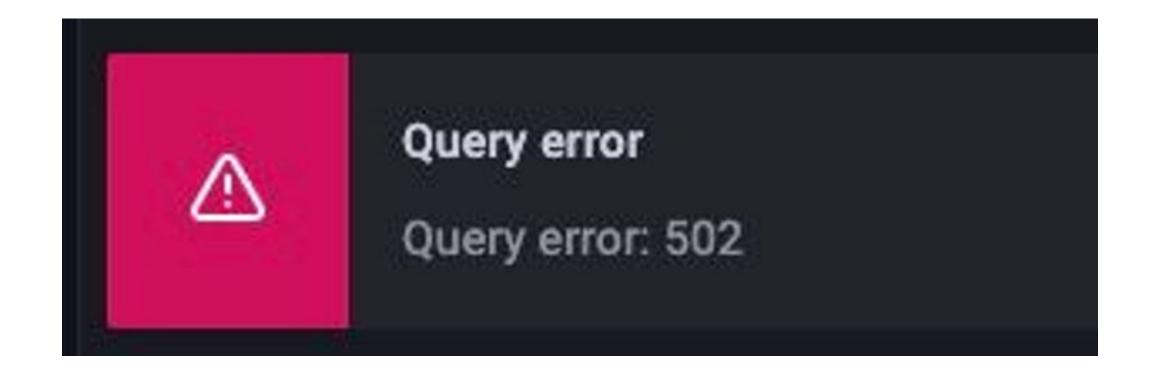
#### DID ME LOOF FOKIS

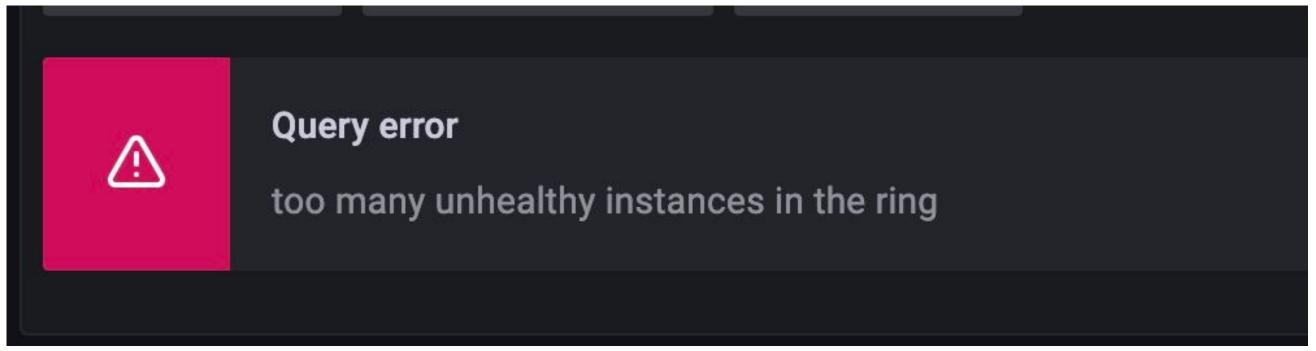


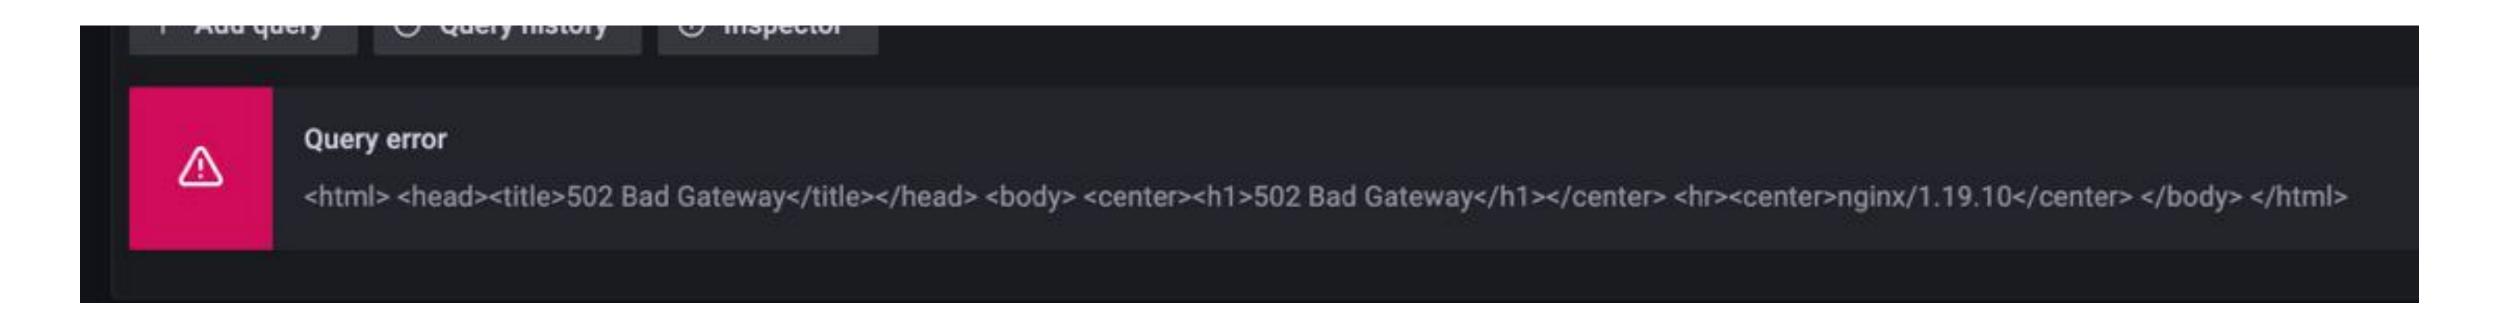
## NOT REALLY

## LOKI WON









#### WHAT IS HAPPENING?



- #4015 Query from Grafana to Loki returning 502 when using 7 days ... Open
- 🕜 #3524 Grafana "bad gateway" / Logcli "EOF", while frontend says "s... 🕑 Closed
- #2540 Querier high memory demands © Closed
- #3753 [loki distributed] slow query (high cpu usage) and time out 
   Closed 
   □

It is time to stop guessing and start thinking

## BLOCK & CHUNK



Ι	ts (varint)	 I	len (uvarint)	 I	log-1 bytes	
I	ts (varint)	 	len (uvarint)	 	log-2 bytes	
1	ts (varint)		len (uvarint)		log-3 bytes	
I 	ts (varint)	 	len (uvarint)	 	log-n bytes	

_	
	MagicNumber(4b)   version(1b)
- 7	block-1 bytes   checksum (4b)
1	block-2 bytes   checksum (4b)
    -	block-n bytes   checksum (4b)
1	#blocks (uvarint)
	<pre>#entries(uvarint)   mint, maxt (varint)   offset, len (uvarint)  </pre>
1	<pre>#entries(uvarint)   mint, maxt (varint)   offset, len (uvarint)  </pre>
  -	<pre>#entries(uvarint)   mint, maxt (varint)   offset, len (uvarint)  </pre>
1	<pre>#entries(uvarint)   mint, maxt (varint)   offset, len (uvarint)  </pre>
	checksum(from #blocks)
  -	metasOffset - offset to the point with #blocks

## BLOCK & CHUNK



	ts (varint)	 	len (uvarint)	 	log-1 bytes	
1	ts (varint)		len (uvarint)		log-2 bytes	1
1	ts (varint)	I	len (uvarint)	I	log-3 bytes	1
I	ts (varint)	I	len (uvarint)	I	log-n bytes	I

1	block-1 bytes   checksum (4b)								
1	block-2 bytes   checksum (4b)								
    -	block-n bytes   checksum (4b)								
	#blocks (uvarint)								
	#entries(uvarint)   mint, maxt (varint)   offset, len (uvarint)								
 	#entries(uvarint)   mint, maxt (varint)   offset, len (uvarint)								
	#entries(uvarint)   mint, maxt (varint)   offset, len (uvarint)								
	#entries(uvarint)   mint, maxt (varint)   offset, len (uvarint)								
	checksum(from #blocks)								
    -	metasOffset - offset to the point with #blocks								

#### INVESTIGATION STAGE



```
for i := 0; i < len(entries); i++ {</pre>
302
                       chunk := &s.chunks[len(s.chunks)-1]
303
304
                       if chunk.closed || !chunk.chunk.SpaceFor(&entries[i]) || s.cutChunkForSynchronization(
                               chunk = s.cutChunk(ctx)
305
306
307
308
                       chunk.lastUpdated = time.Now()
                       if err := chunk.chunk.Append(&entries[i]); err != nil {
309
                               invalid = append(invalid, entryWithError{&entries[i], err})
310
311
                               if chunkenc.IsOutOfOrderErr(err) {
312
                                       outOfOrderSamples++
313
                                       outOfOrderBytes += len(entries[i].Line)
314
315
                               continue
316
```

#### INVESTIGATION STAGE



```
\sim \sim
         // Append implements Chunk.
   667
         func (c *MemChunk) Append(entry *logproto.Entry) error {
                  entryTimestamp := entry.Timestamp.UnixNano()
   669
   670
                 // If the head block is empty but there are cut blocks, we have to make
   671
                 // sure the new entry is not out of order compared to the previous block
   672
                  if c.headFmt < UnorderedHeadBlockFmt && c.head.IsEmpty() && len(c.blocks) > 0 && c.blocks[len(c.blocks)-1].maxt > entryTimestamp {
   673
                          return ErrOutOfOrder
   674
   675
   676
                 if err := c.head.Append(entryTimestamp, entry.Line); err != nil {
••• 677
   678
                          return err
   679
   680
                  if c.head.UncompressedSize() >= c.blockSize {
   681
                          return c.cut()
   682
   683
   684
                  return nil
   685
   686
```

#### INVESTIGATION STAGE



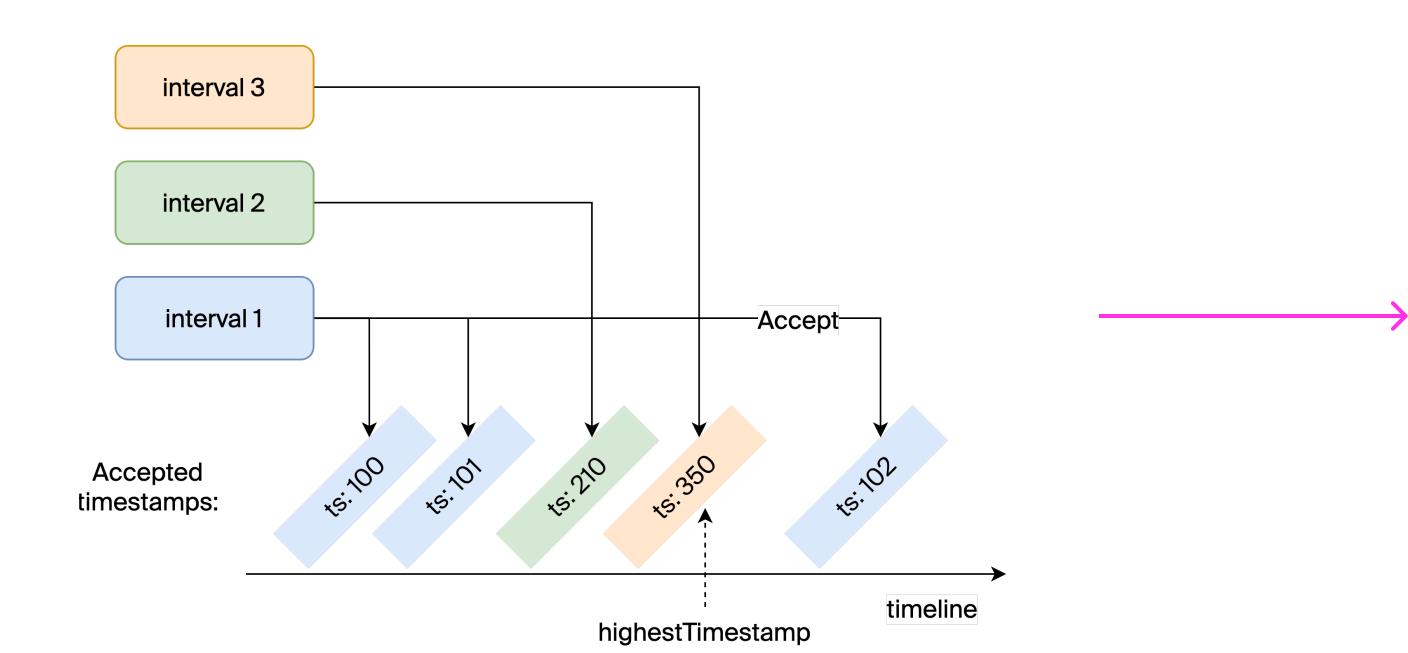
```
It just updates the boundaries
```

```
func (hb *unorderedHeadBlock) Append(ts int64, line string) error {
133
              // Update hb metdata
134
              if hb.size == 0 || hb.mint > ts {
135
                       hb.mint = ts
136
137
              if hb.maxt < ts {</pre>
138
                      hb.maxt = ts
139
140
141
              hb.size += len(line)
142
              hb.lines++
143
144
145
              return nil
146
```

https://github.com/grafana/loki/blob/c75b822fc6998ca57bf53451ec6dc2038c7c1a5e/pkg/chunkenc/unordered.go#L104

## SO MHAT?

## BLOCKS MIX UP





#### **Head Block**

mint	100
maxt	350

ts: 100

ts: 101

ts: 210

ts: 350

ts: 103

#### Chunk

block #1 [100; 350]

block #2 [104; 360]

. . .

block #10 [110; 380]



## CHUNKS MIX UP AS WELL



#### Chunk

block #1 [100; 350]

block #2 [104; 360]

• • •

block #10 [110; 380]

#### Chunk

block #1 [112; 381]

block #2 [115; 390]

. . .

block #10 [116; 720]

#### Chunk

block #1 [117; 200]

block #2 [260; 500]

. . .

block #10 [120; 800]



#### Chunk

block #1 [100; 350]

block #2 [104; 360]

. . .

block #10 [110; 380]

#### Chunk

block #1 [112; 381]

block #2 [115; 390]

•••

block #10 [116; 720]

#### Chunk

block #1 [117; 200]

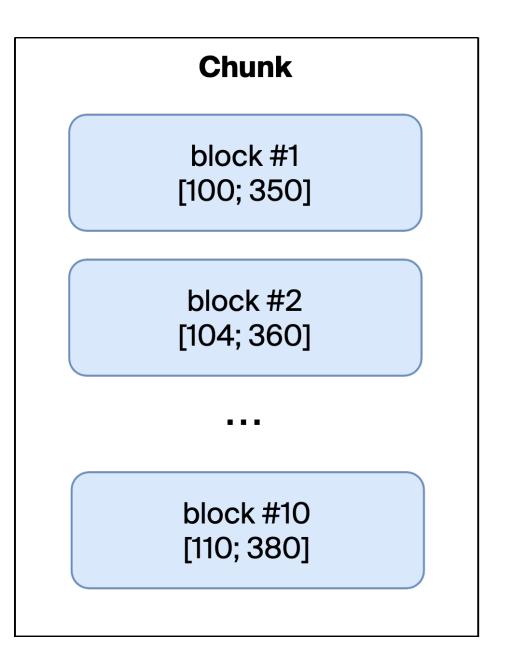
block #2 [260; 500]

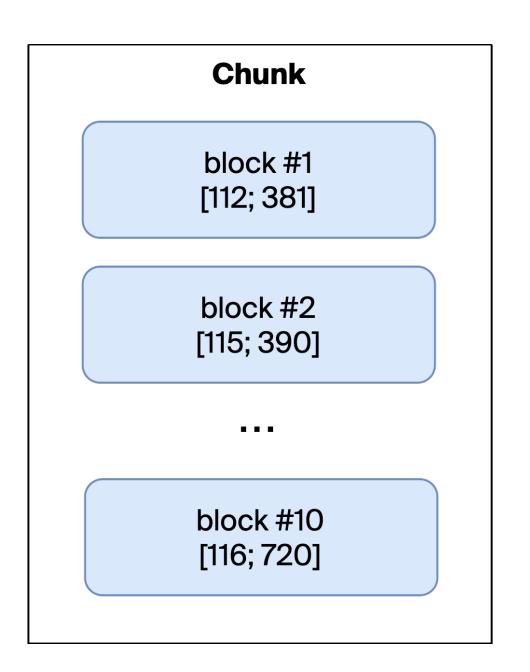
. . .

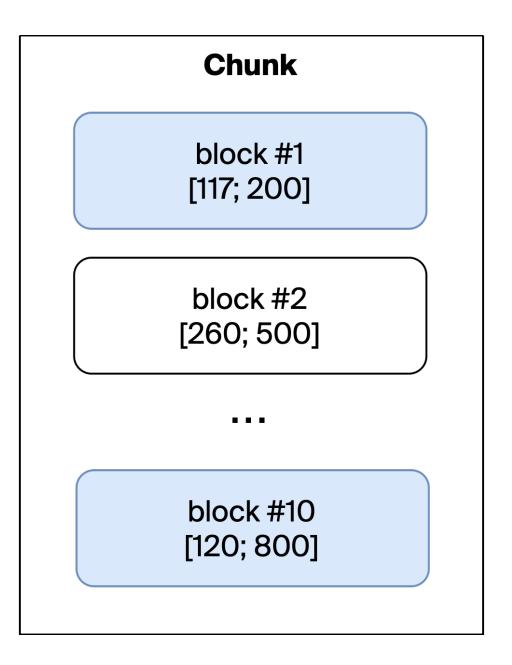
block #10 [120; 800]



- Loki is forced to download almost all chunks from the storage
- Process in memory
- Resort & filter



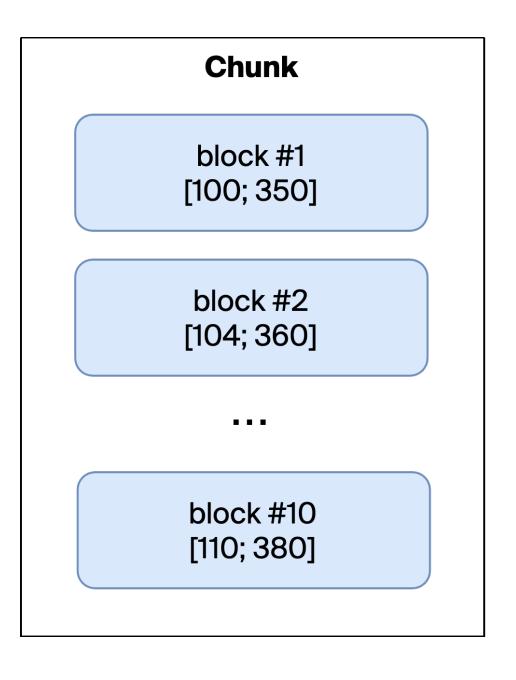


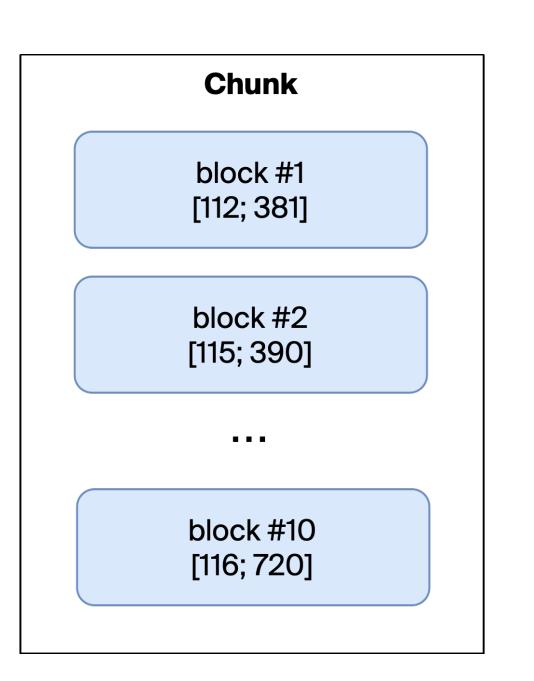


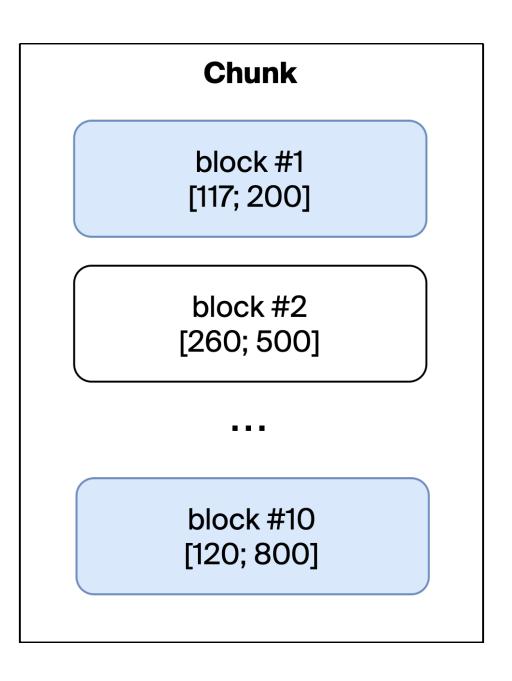


- Loki is forced to download almost all chunks from the storage
- Process in memory
- Resort & filter









## HOW TO FIX EVERYTHING?



Logs		
Labels	Message	
<pre>{component="printer", location="f2c16", level="error"}</pre>	"Out of paper"	chunk #1 stream 1
<pre>{component="printer", location="f2c16", level="error"}</pre>	"Too much paper"	Sugaini
<pre>{component="supplier", location="f2c16", level="info"}</pre>	"Paper exhasted"	chunk #2 stream 2

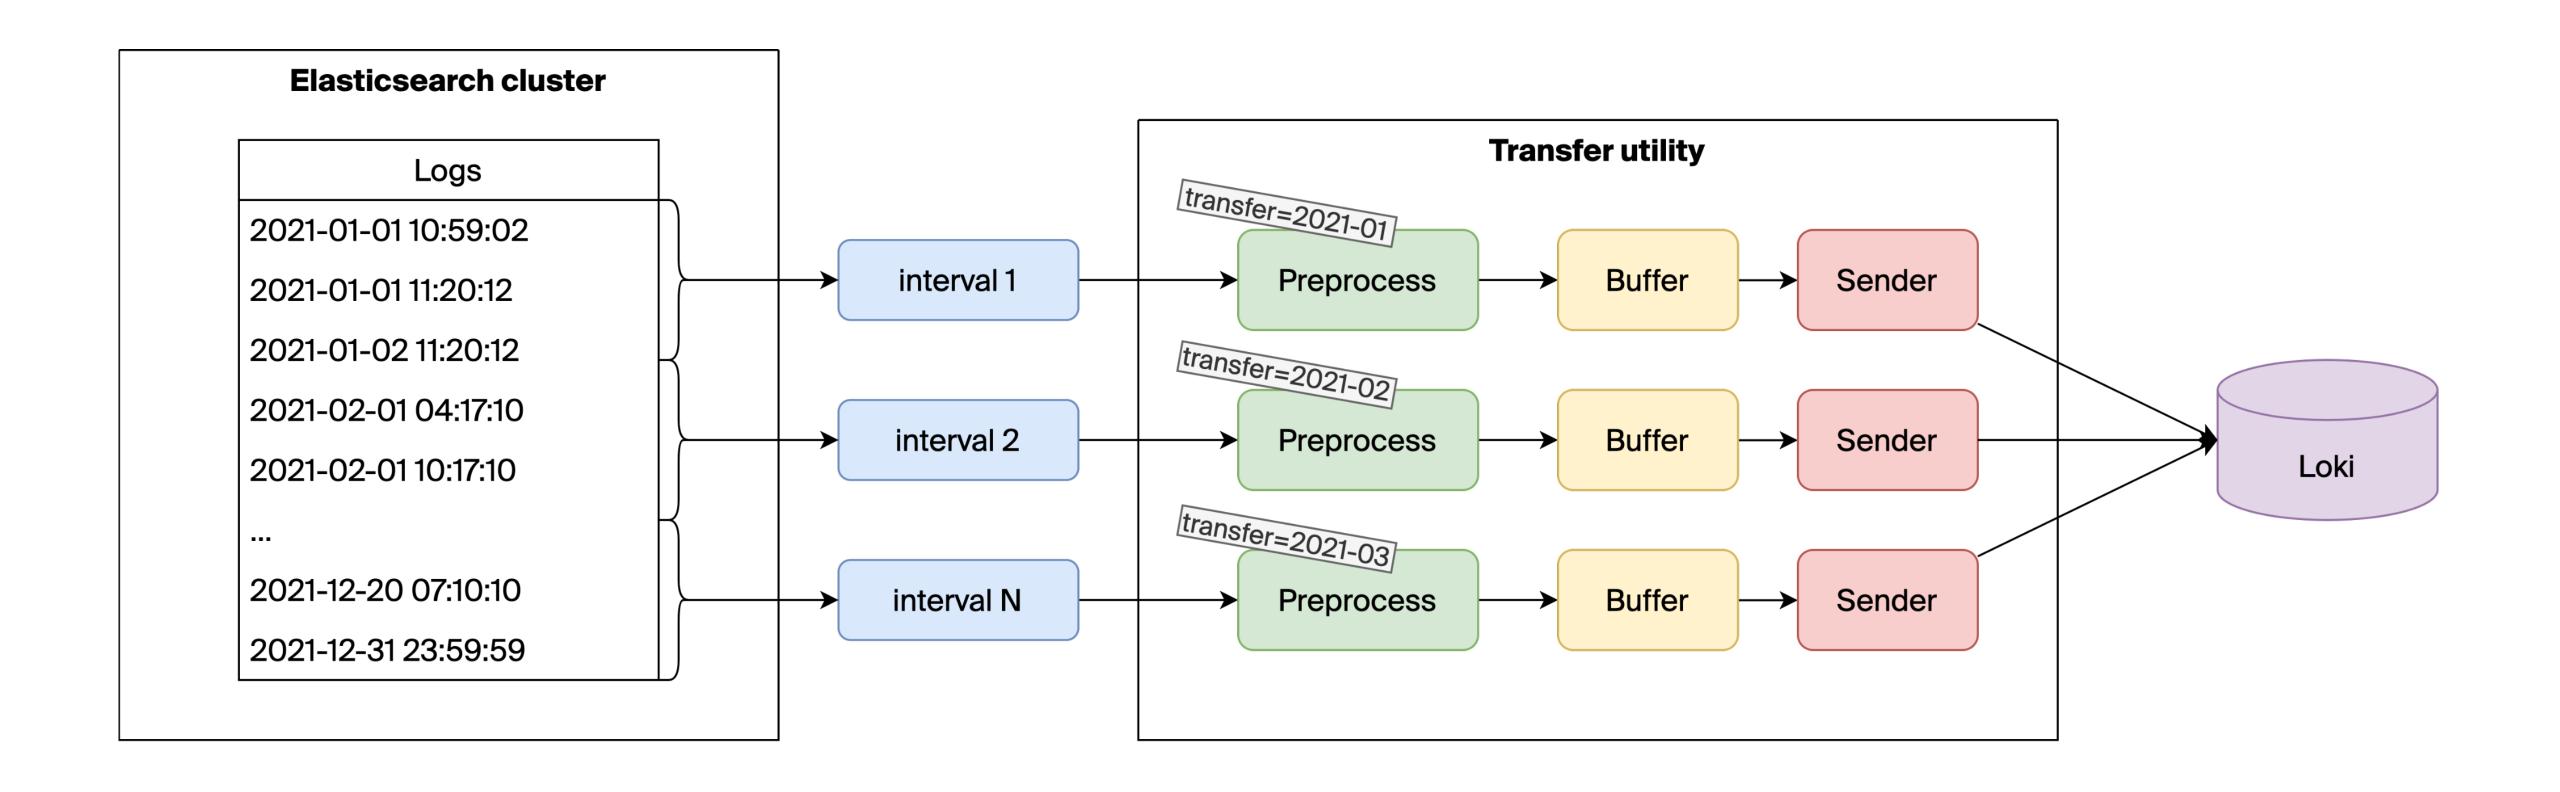
#### HOW TO FIX EVERYTHING?



- 1. Don't set max\_chunk\_age to a high value
- 2. Don't send logs in parallel for the same set of labels (i.e. a stream)
- 3. You may create additional streams by adding new labels
- 4. Send logs within a stream only in the strict order

#### JUST ADD LABELS







#### Chunk

block #1 [100; 120]

block #2 [121; 130]

• • •

block #10 [131; 140]

#### Chunk

block #1 [141; 160]

block #2 [161; 200]

• • •

block #10 [201; 220]

#### Chunk

block #1 [221; 250]

block #2 [251; 270]

. . .

block #10 [291; 300]

## CHUNKS SITUATION



Query matches less chunks



2 All blocks within all chunks are strictly ordered

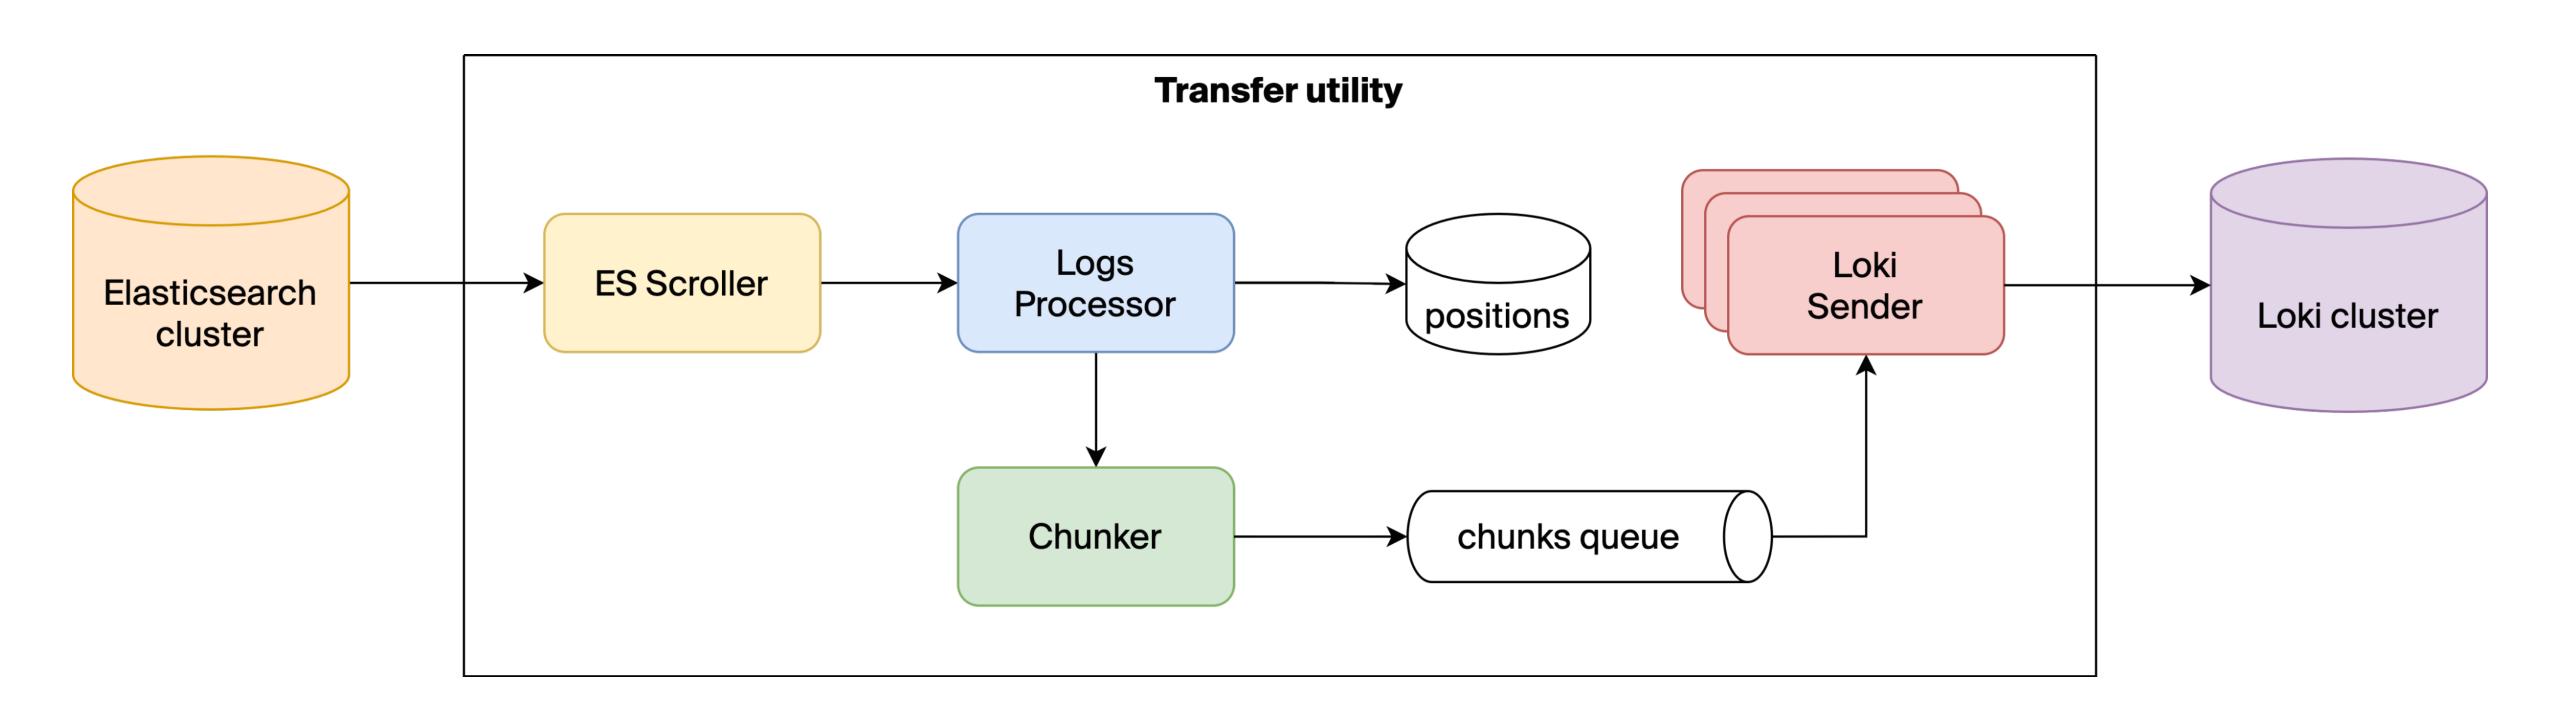


3 Fast and efficient query execution

## TRANSFER IMPLEMENTATION







#### ES SCROLLER



- Sort by (@timestamp, log.offset)
- Use search\_after to scroll & resume search
- 3 Do not let a consumer wait for data
- 4 Handle errors properly
- Some shards may return no data

ALWAYS WAIT FOR ERROR RESOLUTION - WE CANNOT LOOSE LOG LINES

#### LOGS PROCESSOR



PERFORMS A SIMILAR PROCESSING AS PROMTAIL

- Extract timestamp
- 2 Extract labels
- 3 Filter logs
- Make a unique stream to ensure parallelism

#### LOGS PROCESSOR



#### LOKI SENDER

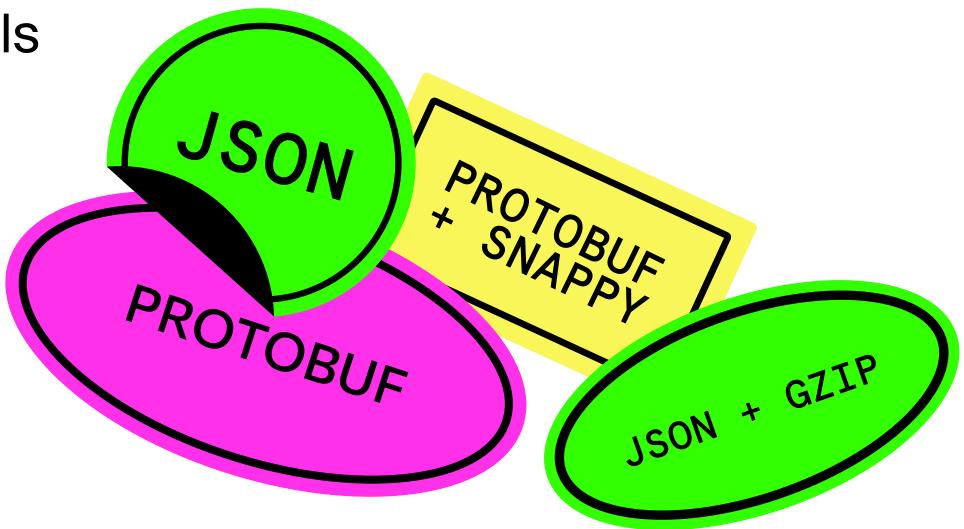


- Send data to <u>/loki/api/v1/push</u>
- 2 Send data sequentially
- 3 Use only 1 worker (for now)
- Support for different protocols

#### LOKI SENDER



- Send data to /loki/api/v1/push
- 2 Send data sequentially
- 3 Use only 1 worker (for now)
- 4 Support for different protocols



## YOU CAN ALREADY USE IT





HTTPS://GITHUB.COM/KTSSTUDIO/ES2LOKI

#### YOU CAN ALREADY USE IT



```
1 class Transfer(BaseTransfer):
      def extract_doc_labels(self, source: dict) → Optional[MutableMapping[str, str]]:
          return dict(
              app=source.get("fields", {}).get("service_name"),
              job="logs",
              level=source.get("level"),
              node_name=source.get("host", {}).get("name"),
              logger_name=source.get("logger_name"),
                            ELASTIC_HOSTS=http://localhost:9200 \
```

ELASTIC\_INDEX="filebeat-\*" \

python ./transfer.py

LOKI\_URL=http://localhost:3100 \

#### YOU CAN ALREADY USE IT



JUST ADD YOUR OWN DOCKER IMAGE, WE'LL DO THE REST:

```
$ helm repo add kts https://charts.kts.studio
$ helm repo update
$ helm upgrade --install \
   RELEASE_NAME \
   kts/es2loki \
   --set image.repository=your-docker-image
   --set image.tag=latest
```

## SEE HOW IT WORKS

#### Components

- 1. Elasticsearch
- 2. Kibana
- 3. filebeat (imports "old" logs to Elasticsearch)
- 4. Grafana
- 5. Loki
- 6. Promtail (imports "new" logs to Loki)
- 7. PostgreSQL (needed for es2loki)
- 8. es2loki

#### Usage

In order to run a demo you may use:

docker compose up



HTTPS://GITHUB.COM/KTSSTUDIO/ES2LOKI

## ANY PROFIT?



## LOKI IN KUBERNETES



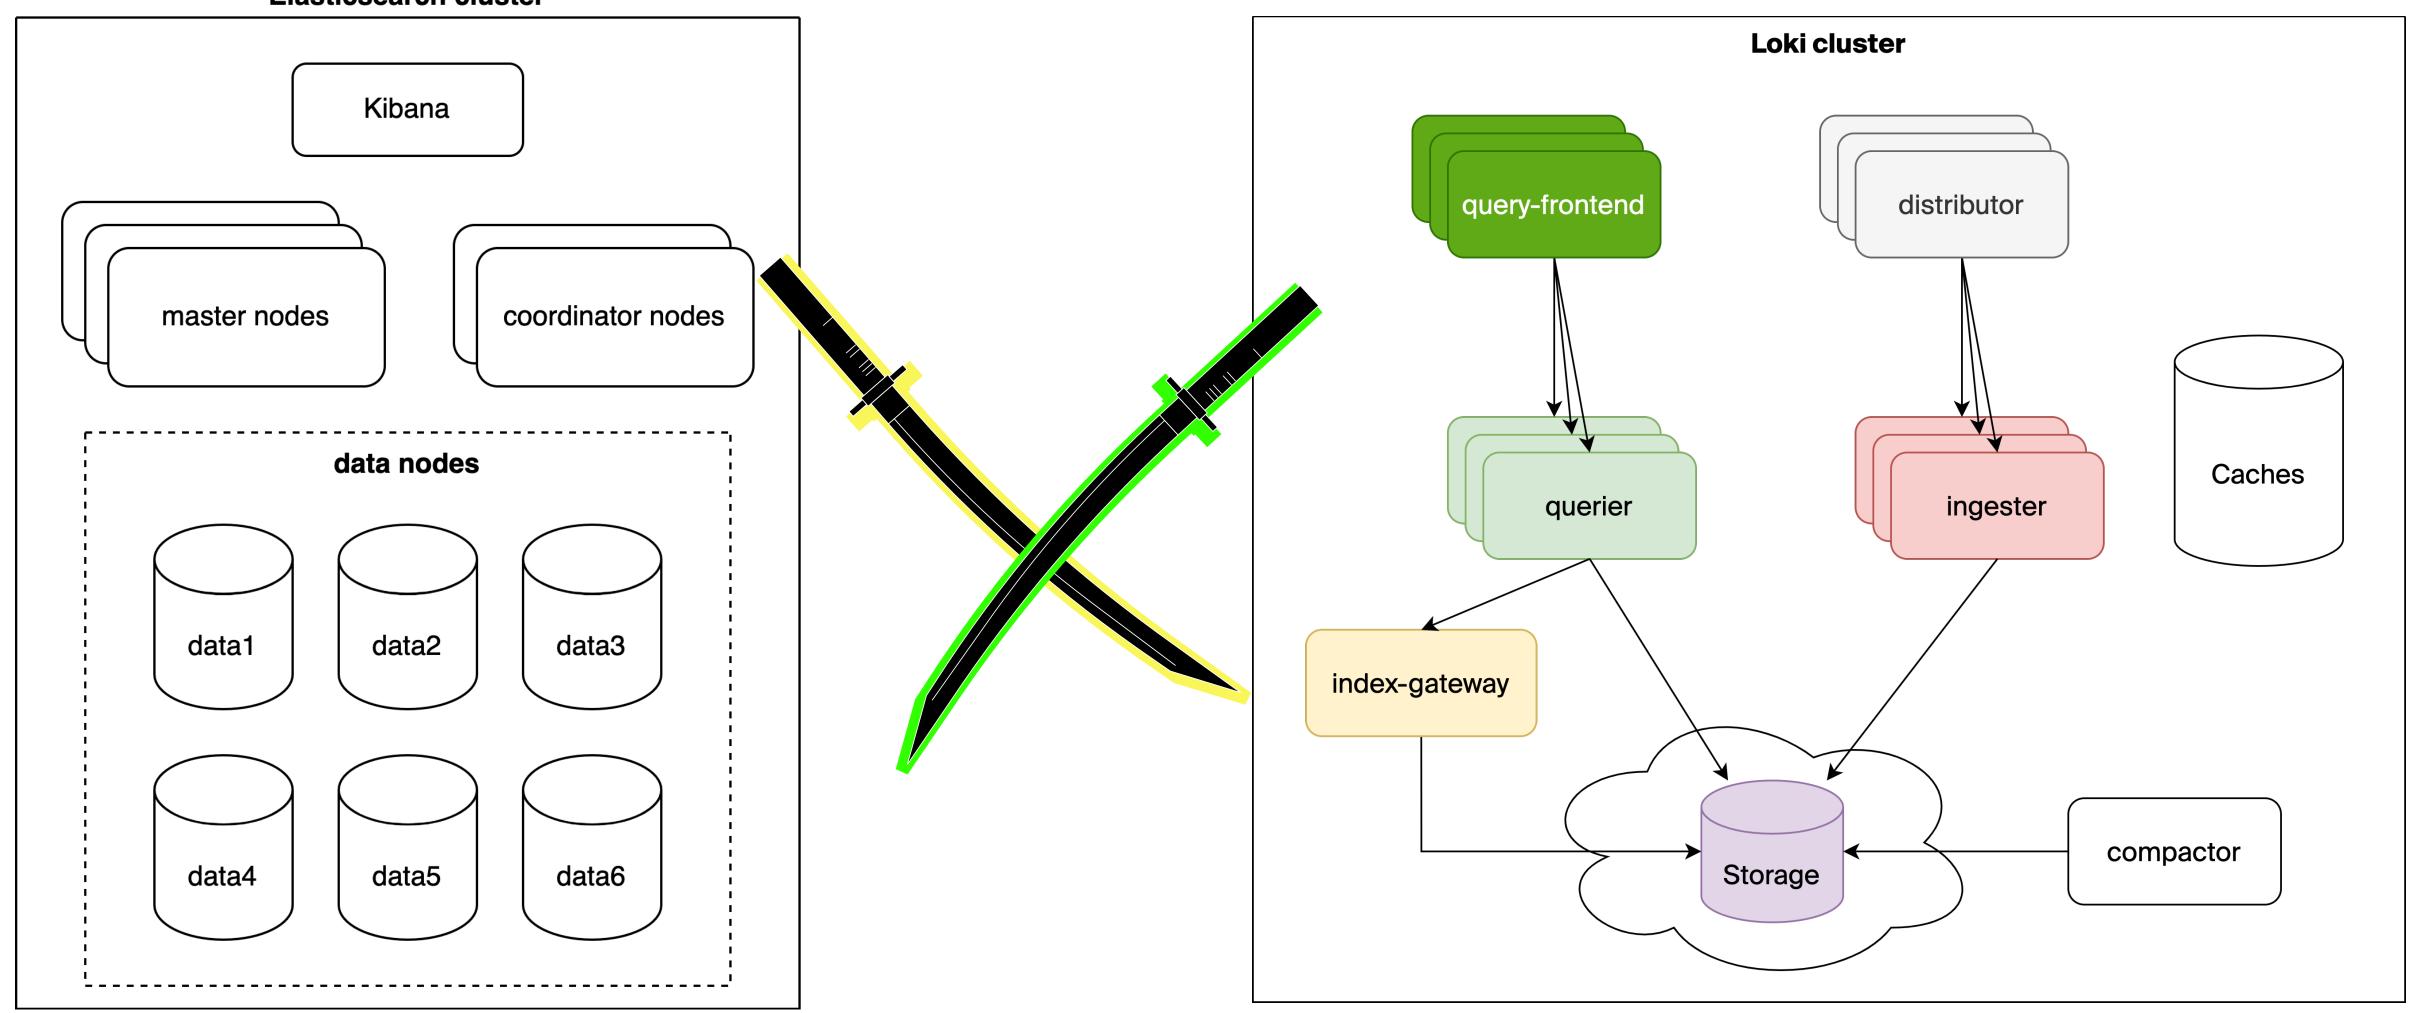
		naci	Kelneal 3a Loll	— Po	ds(lol	(i)[19] -								
NAME↑	PF	READY	RESTARTS STATUS	CPU	MEM	CPU/R:L	MEM/R:L	%CPU/R	%CPU/L	%MEM/R	%MEM/L	IP	NODE	Q
loki-compactor-565b5b8565-wczqw	•	1/1	1 Running	16	58	30:0	100:600	53	n/a	58	9	10.233.88.135	wnode1	В
loki-distributor-64b94c49fc-8kpxl	•	1/1	0 Running	30	77	100:0	100:2048	30	n/a	77	3	10.233.88.30	wnode1	E
loki-distributor-64b94c49fc-ffq2z	•	1/1	2 Running	23	73	100:0	100:2048	23	n/a	73	3	10.233.73.103	wnode2	Е
loki-distributor-64b94c49fc-q9ckb	•	1/1	0 Running	26	71	100:0	100:2048	26	n/a	71	3	10.233.126.26	wnode3	E
loki-gateway-55dcf67678-w8qww	•	1/1	0 Running	701	14	1000:0	50:500	70	n/a	28	2	10.233.73.134	wnode2	В
loki-index-gateway-0	•	1/1	2 Running	2	33	30:0	50:200	6	n/a	66	16	10.233.88.22	wnode1	В
loki-ingester-0	•	1/1	0 Running	32	899	100:0	3072:7168	32	n/a	29	12	10.233.88.207	wnode1	В
loki-ingester-1	•	1/1	0 Running	30	1233	100:0	3072:7168	30	n/a	40	17	10.233.88.185	wnode1	В
loki-ingester-2	•	1/1	48 Running	31	1459	100:0	3072:7168	31	n/a	47	20	10.233.88.216	wnode1	В
loki-memcached-frontend-0	•	2/2	0 Running	1	31	50:0	50:0	2	n/a	62	n/a	10.233.88.245	wnode1	E
loki-memcached-frontend-1	•	2/2	0 Running	2	31	50:0	50:0	4	n/a	63	n/a	10.233.126.24	wnode3	В
loki-memcached-frontend-2	•	2/2	1 Running	1	26	50:0	50:0	2	n/a	53	n/a	10.233.73.88	wnode2	В
loki-memcached-index-queries-0	• 0	2/2	0 Running	3	58	10:0	100:0	30	n/a	58	n/a	10.233.88.26	wnode1	В
loki-memcached-index-queries-1	•	2/2	0 Running	2	59	10:0	100:0	20	n/a	59	n/a	10.233.126.92	wnode3	В
loki-memcached-index-queries-2	•	2/2	2 Running	1	56	10:0	100:0	10	n/a	56	n/a	10.233.73.86	wnode2	В
loki-querier-866486df4d-2lhp7	•	1/1	0 Running	11	439	100:0	1024:5120	11	n/a	42	8	10.233.126.18	wnode3	В
loki-querier-866486df4d-lw94f	• 7	1/1	4 Running	8	453	100:0	1024:5120	8	n/a	44	8	10.233.73.101	wnode2	E
loki-querier-866486df4d-rmzw8	•	1/1	0 Running	12	441	100:0	1024:5120	12	n/a	43	8	10.233.88.58	wnode1	В
loki-query-frontend-7f89d8c6c7-5wn9j	• 11	1/1	0 Running	HCT1	109	100:0	300:5120	1	n/a	36	2	10.233.126.219	wnode3	E

<namespace> <pod>

## ELASTICSEARCH VS LOKI

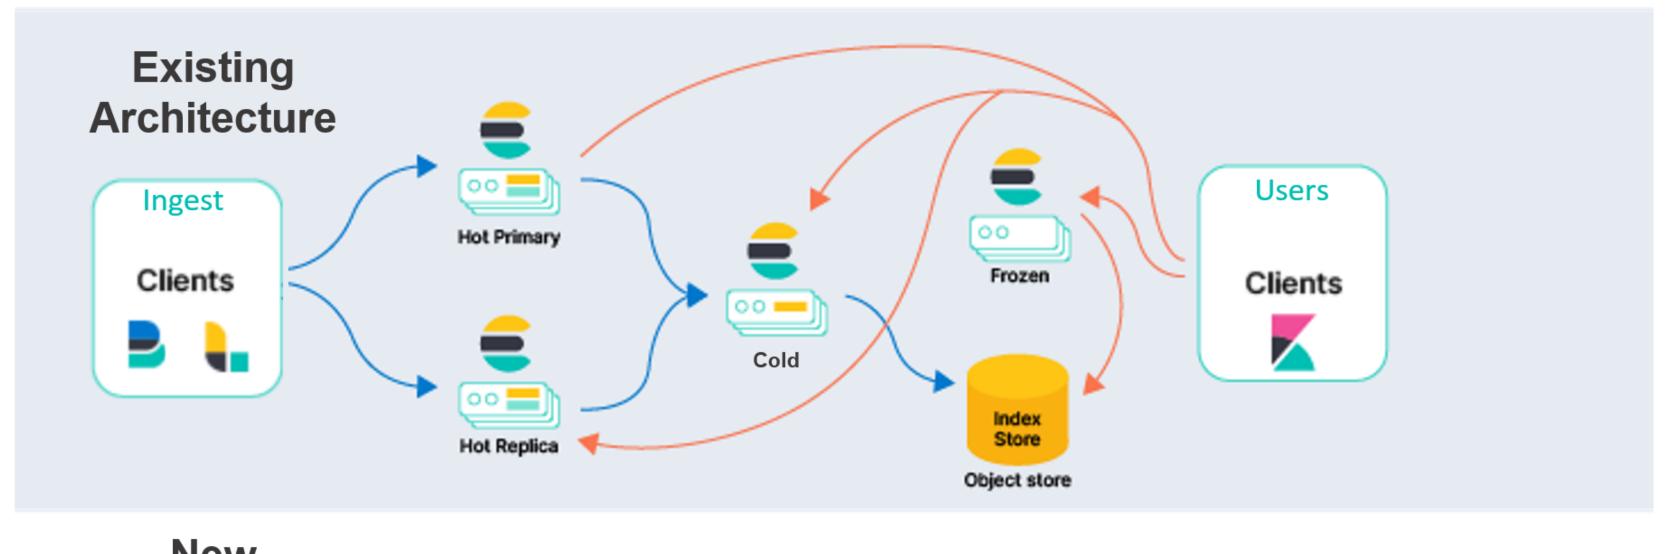


#### Elasticsearch cluster



## NEW STATELESS ELASTICSEARCH ARCHITECTURE







### INFRA COMPARISON



#### **ELASTICSEARCH**

→ CPU: 44

→ Memory: 232 GB

→ Main disks: 338 GB

→ Storage disks: 25 TB



 $\rightarrow$  CPU: 7

→ Memory: 14 GB

→ Main disks: 105 GB

→ S3 Storage: 5 TB



## INFRA COMPARISON



#### **ELASTICSEARCH**

→ CPU: 44

Memory: 232 GB

Main disks: 338 GB

Storage disks: 25 TB

#### LOKI

→ CPU: 7

→ Memory: 14 GB

→ Main disks: 105 GB

→ S3 Storage: 5 TB

#### **SAVINGS**





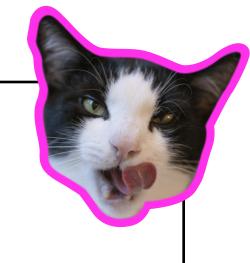
→ 16x less RAM

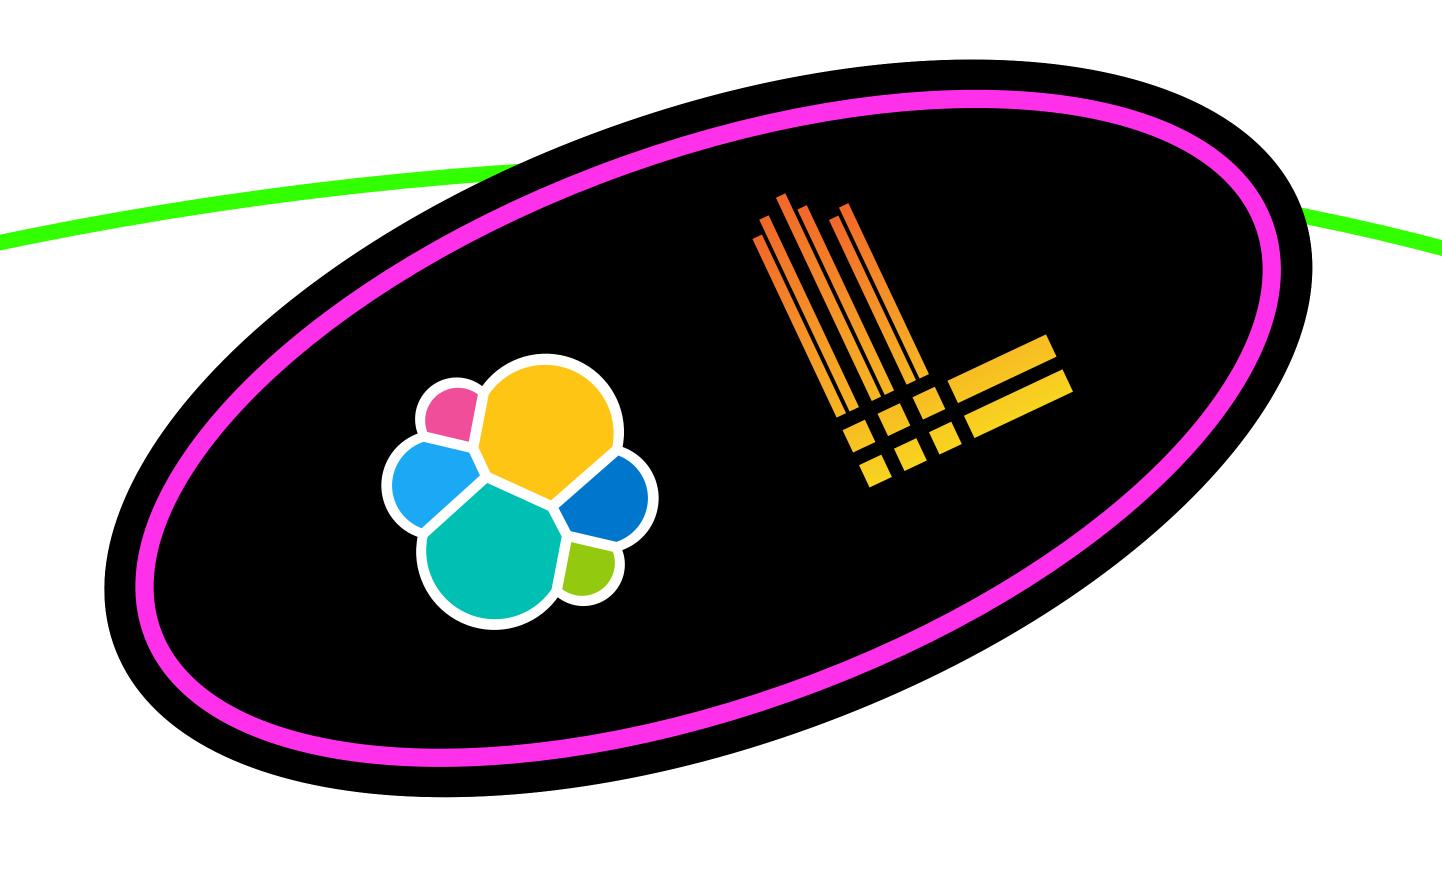
→ 3x less disk

→ 5x less main storage









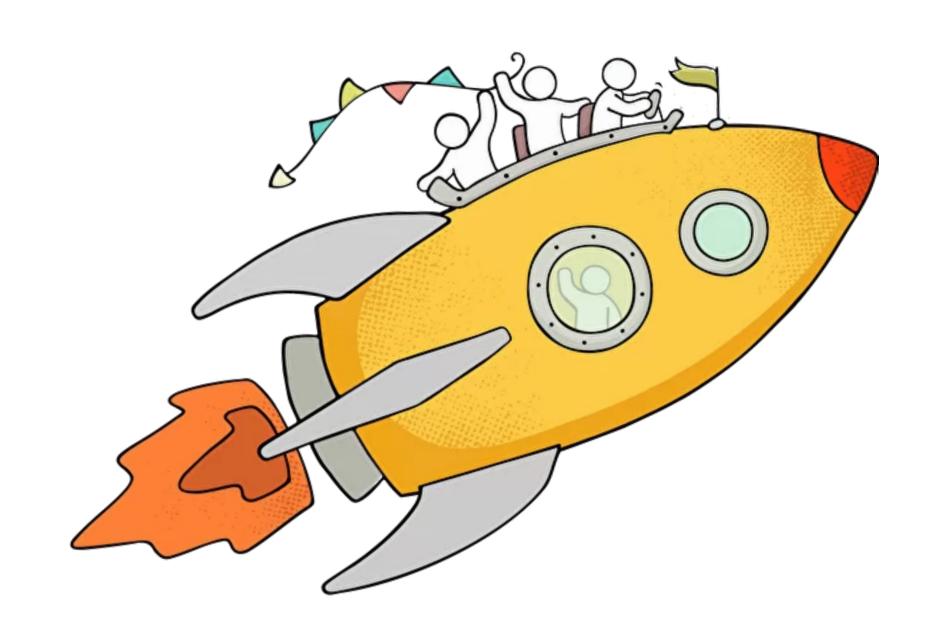


## RESULTS



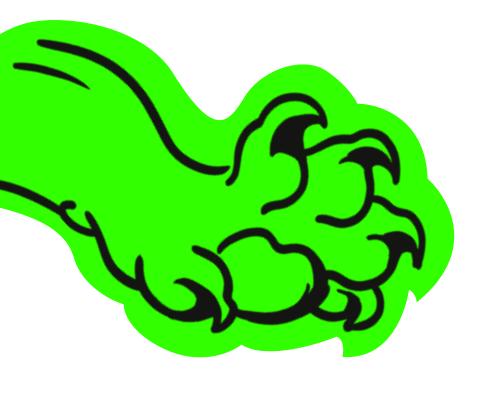
#### POSITIVE:

- 4x logs infrastructure cost reduction
- 2 Simplier infrastructure maintanance
- Much more stable installation than Elasticsearch
- Logs transfer in adequate timespan
- 5 Easier to scale



## RESULTS

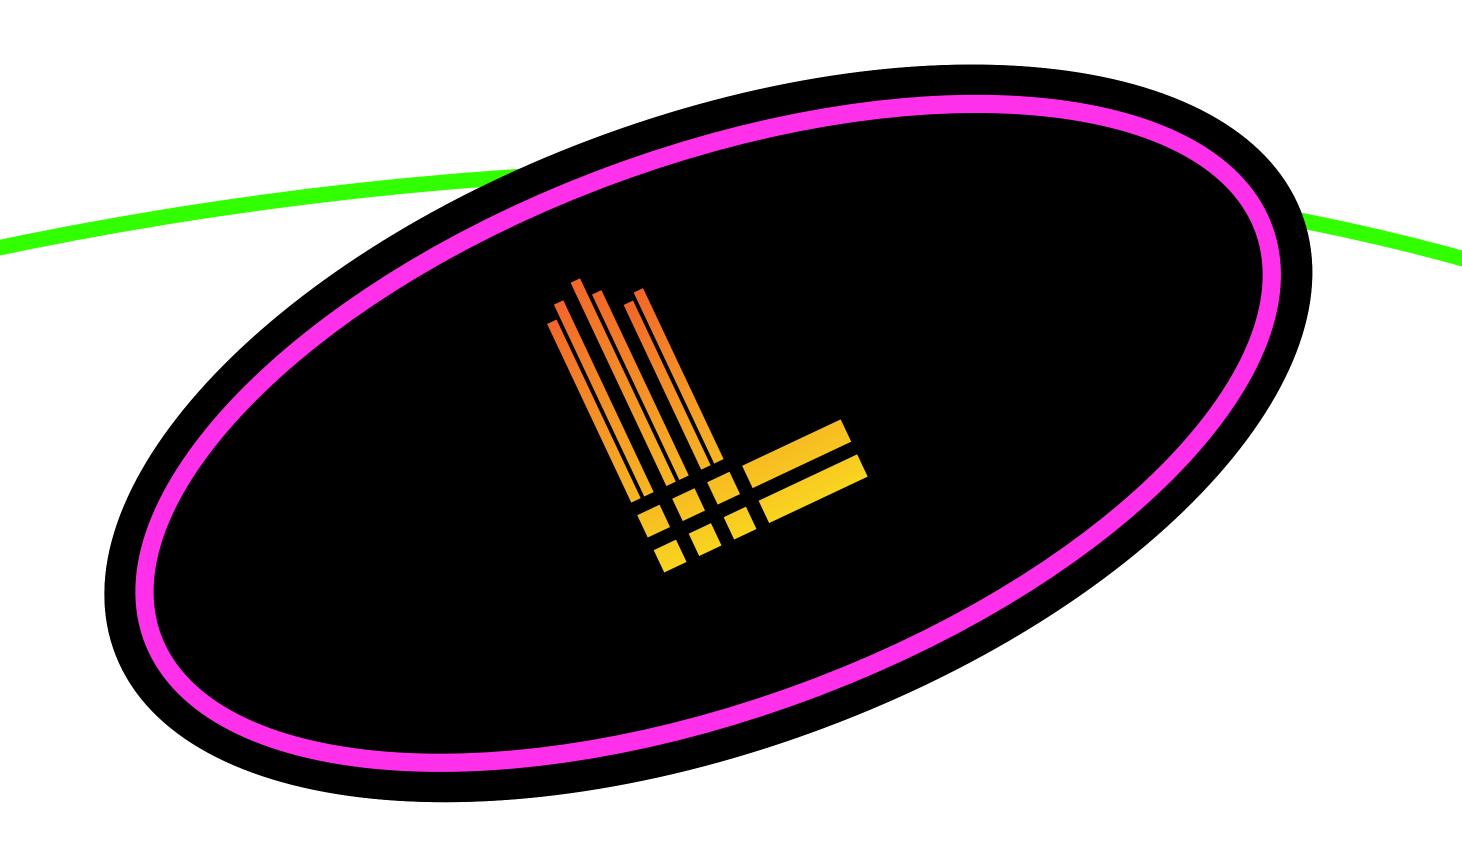






#### NOT SO POSITIVE, BUT OK:

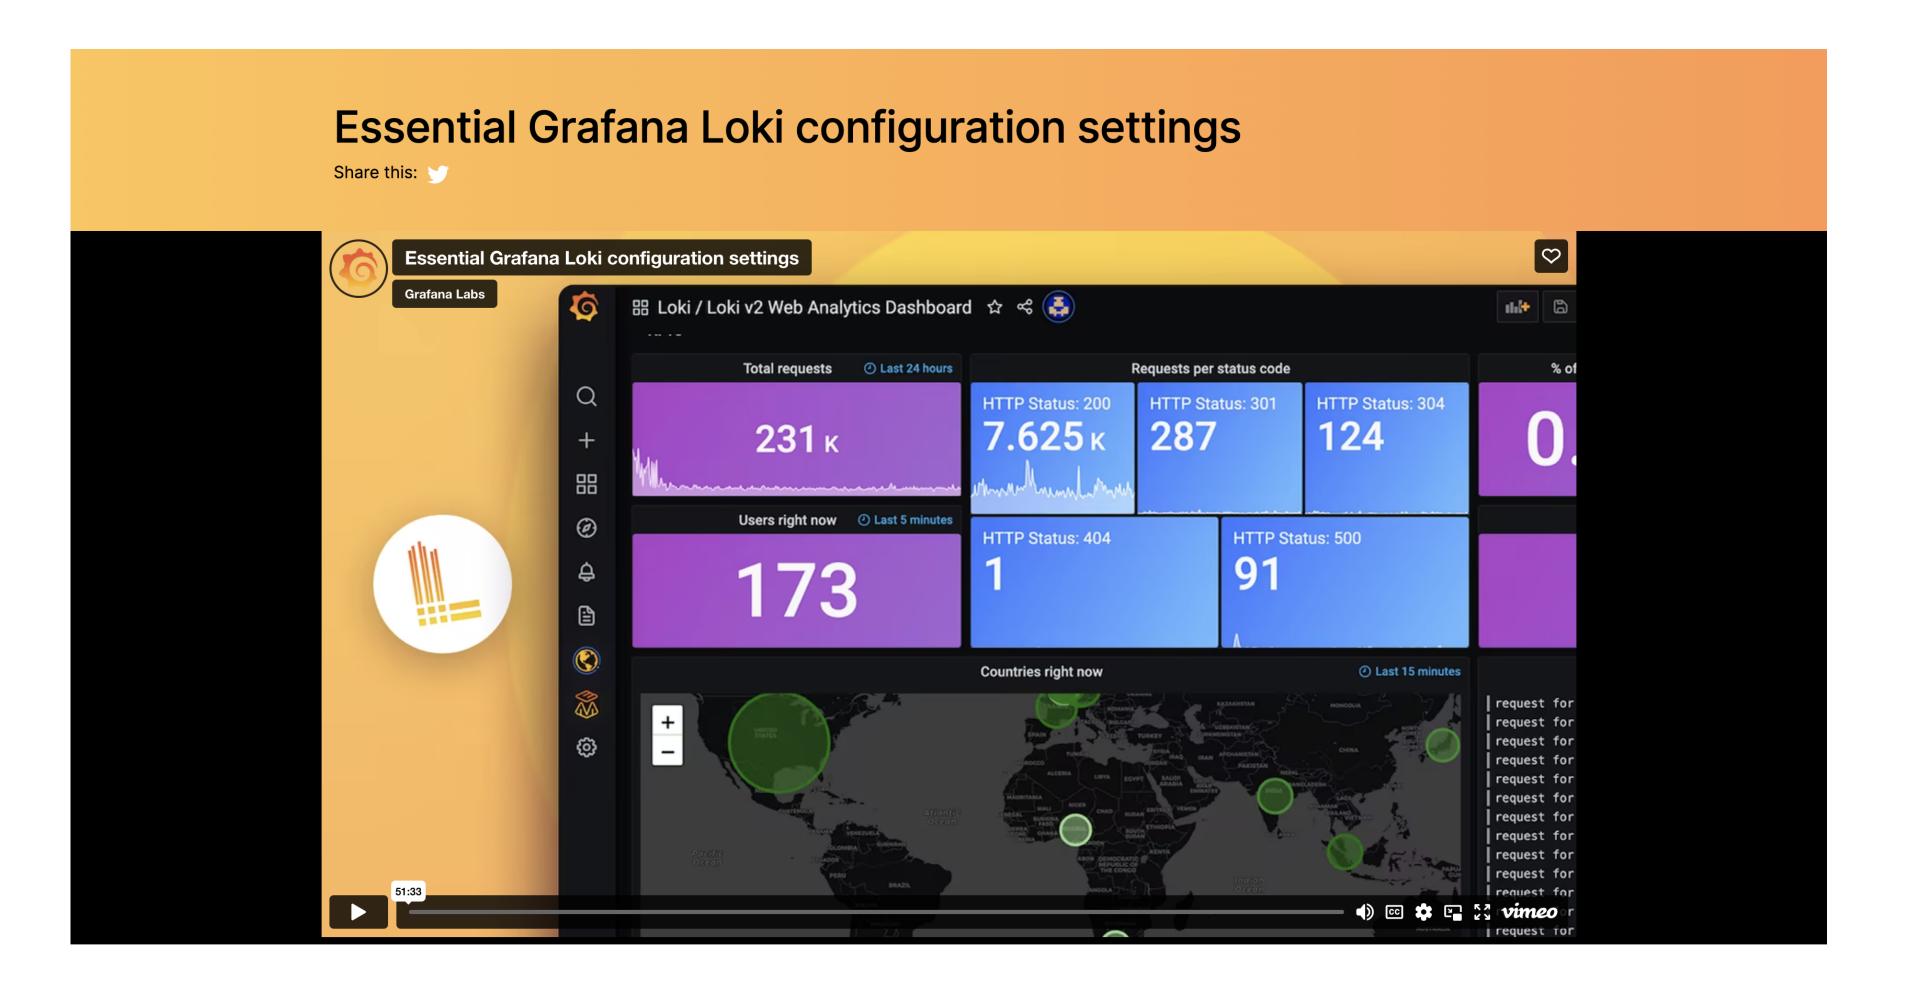
- Needed to use faster disks for S3
- Grafana is not Kibana
- 3 No full-text search
- No Machine Learning
- Some queries in Loki are slower



# LOKI CONFIGURATION --> CHECKLIST (BONUS)

#### CONFIGURATION TIPS







## TECH TALKS - DECEMBER, 16



MLOPS
IN ENTERPRISES
AND SMB

→ 15:30 - 15:50



#### Leave your feedback!

You can rate the talk and give a feedback on what you've liked or what could be improved

